

2012-12-14

Development of the Finite Difference Time Domain Method on a Lebedev Grid for Anisotropic Materials

Nauta, Marcel

Nauta, M. (2012). Development of the Finite Difference Time Domain Method on a Lebedev Grid for Anisotropic Materials (Master's thesis, University of Calgary, Calgary, Canada).

Retrieved from <https://prism.ucalgary.ca>. doi:10.11575/PRISM/25689

<http://hdl.handle.net/11023/355>

Downloaded from PRISM Repository, University of Calgary

UNIVERSITY OF CALGARY

Development of the Finite Difference Time Domain Method on a Lebedev Grid for
Anisotropic Materials

by

Marcel D. Nauta

A THESIS

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE
DEGREE OF MASTER OF SCIENCE

DEPARTMENT OF ELECTRICAL AND COMPUTER ENGINEERING

CALGARY, ALBERTA

December, 2012

© Marcel D. Nauta 2012

Abstract

The finite-difference time-domain (FDTD) method is derived on a Lebedev grid, instead of the standard Yee grid, to better represent the constitutive relations in anisotropic materials. The Lebedev grid extends the Yee grid by approximating Maxwell's equations with tensor constitutive relations using only central differences. A dispersion relation with stability criteria is derived and it is proven that the Lebedev grid has a consistent calculus. An integral derivation of the update equations illustrates how to appropriately excite the grid. This approach is also used to derive the update equations at planar material interfaces and domain edge PEC. Lebedev grid results are compared with analytical and Yee grid solutions using an equal memory comparison. Numerical results show that the Lebedev grid suffers greater dispersion error but better represents material interfaces. Focus is given to generalizing the concepts that make the Yee grid robust for isotropic materials.

Keywords: FDTD, anisotropic materials, Lebedev grid, collocated grids

Acknowledgements

There are many friends and colleagues that have helped me throughout the course of my education, but a few stand out from the crowd.

Thomas Apperley has been my fellow physicist working in engineering and has been an amazing help in all of the little difficulties that come up in the transition between the two fields. He has always been quick to give me honest advice on my writing style, figures, and the way information is presented. I am thankful for his patience with my questions and the time he spent reading through this thesis.

I must also thank my supervisors for their advice throughout my Masters, Mike Potter and Michal Okoniewski. Most of the novel ideas presented in this thesis have come out of lengthy discussions on little details in my derivations. They have also spoiled me and sent me to several conferences that allowed me to discuss approaches with other researchers in the field.

Finally, I would like to thank Billy Wu. He is the unthanked hard worker in our group that deals with little details and pays attention to the welfare of fellow group members. His concern and diligence is appreciated.

Table of Contents

Abstract	ii
Acknowledgements	iii
Table of Contents	iv
List of Tables	vi
List of Figures	vii
List of Abbreviations and Symbols	xiii
1 Introduction	1
1.1 Motivation	1
1.2 Thesis Goals	2
1.3 Thesis Outline	3
2 Background	5
2.1 Governing Equations	6
2.2 Electromagnetic Effects of Anisotropic Dielectrics	10
2.3 Sapphire Permittivity Tensor	16
2.4 FDTD Techniques for Handling Anisotropy	17
3 Lebedev Grid in an Unbounded Domain	20
3.1 Notation	21
3.2 Central Difference Based Derivation of the Update Equations	22
3.3 Grid Structure	25
3.4 Dispersion and Stability Analysis	26
3.5 Equal Memory Comparison	33
3.6 Grid Degeneracy	35
3.7 Integral Derivation of Update Equations	36
3.8 Sources and Sampling	40
3.9 Plane Waves in an Anisotropic Material	45
3.10 Conservation of Charge	50
3.11 Infinite Domain Summary	53
4 Planar Boundary Conditions and Material Interfaces	55
4.1 Convolutional Perfectly Matched Layers	55

4.1.1	Implementation	56
4.1.2	Measured Reflection Coefficient	58
4.2	Perfect Electric Conductors	59
4.2.1	Derivation	61
4.2.2	Implementation	63
4.2.3	Analysis of a Filled Rectangular Waveguide	66
4.3	Material Interfaces	70
4.3.1	Derivation	71
4.3.2	Implementation	74
4.3.3	Reflection and Transmission of Normally Incident Plane Waves	75
4.4	Planar Boundary Conditions Summary	81
5	Grid Extensions	83
5.1	Ferrites Simulated by Gilbert's Equation of Motion	83
5.2	Anisotropic Dielectrics With an Axis of Symmetry	86
5.3	Grid Extensions Summary	88
6	Conclusions	89
6.1	Accomplishments	90
6.2	Future Work	93
	Bibliography	94
A	Waveforms	100
A.1	Time Domain Waveforms	100
A.2	Spatial Extent of Excitations	104
B	Extracting Data	108
B.1	Isolating Propagating Modes	108
B.1.1	Sources of Error	109
B.2	Transmission Line Cross Sections of Propagating Modes	110
B.3	Measuring Reflection and Transmission	111
C	Programming Structure	113

List of Tables

3.1	Stability limit testing for lossless anisotropic materials in an infinite domain. Each simulation was run in a $7 \times 9 \times 11$ domain with a cell size of $1.3 \times 2.2 \times 1.5$ mm and periodic boundary conditions. A simulation is considered stable if it runs for 65536 timesteps and none of the fields within the domain have a magnitude exceeding 10^{30} . Each material was simulated with its principal axes aligned with the grid axes, rotated by 36° about the z -axis, and further rotated by 11° about the y -axis. The analytical stability limit was calculated using (3.26). The simulated limit was found by running a simulation at $S = 1.05$ and then decreasing it by 0.01 until the simulation was stable. Note that several simulations were measured as stable for $S = 1.01$. Further investigation showed that these simulations were unstable but the exponential increase was too slow to reach the assigned threshold within the simulation time.	30
-----	---	----

List of Figures

2.1	Plane wave velocities as a function of direction for an isotropic dielectric. The phase velocity is given by a vector between the origin and a point on the surface for each polarization. In isotropic media, any polarization has the same phase velocity that is independent of direction.	14
2.2	Plane wave velocities of the two basis polarizations as a function of direction for a uniaxial dielectric. The phase velocity is given by a vector between the origin and a point on the surface for each polarization. In a uniaxial dielectric, there are two basis linear polarization states in any direction of propagation. The phase velocity of one state is independent of direction (ordinary wave) whereas the other depends on direction (extraordinary wave). There is a single axis (optical axis) where the phase velocities for both states are the same. Note that the anisotropy is exaggerated for the sake of illustration.	14
2.3	Plane wave velocities of the two basis polarizations as a function of direction for a biaxial dielectric. The phase velocity is given by a vector between the origin and a point on the surface for each polarization. In a biaxial dielectric, there are two basis linear polarization states in any direction of propagation and their phase velocities depend on direction. There are two axes (optical axes) where the phase velocities for both states are the same. Note that the anisotropy is exaggerated for the sake of illustration.	15
2.4	Plane wave velocities of the two basis polarizations as a function of direction for a gyrotropic dielectric. The phase velocity is given by a vector between the origin and a point on the surface for each polarization. In a gyrotropic dielectric, there are two basis elliptical polarization states in any direction of propagation. The polarization states are of different handedness for opposite directions of propagation. This leads to non-reciprocal behavior.	15
3.1	Indexing notation for the Lebedev grid. These fields are indexed with respect to the electric field located at $(i, j, k) = (2, 1, 3)$. Together the electric and magnetic fields form a Cartesian grid which allows the fields to be stored together in a regular 3D storage format.	26

3.2	Traditional Yee grid (left) and proposed Lebedev grid (right). The Lebedev grid uses collocated field components and staggered field types so that central derivatives can be used to compute the spatial derivative of any field component along any axis. The Yee grid uses staggered field components and staggered field types because it is specialized to compute only the spatial derivatives required for Maxwell's equations in isotropic materials.	27
3.3	Normalized numerical phase velocity error as a function of direction on a Lebedev grid for an isotropic dielectric. The amount of error depends on the number of grid points per wavelength causing dispersion that is worst along the cardinal axes.	32
3.4	Normalized numerical phase velocity error as a function of direction on a Lebedev grid for uniaxial dielectric. The error depends only slightly on the polarization and the magnitude of the error is comparable to isotropic materials.	32
3.5	Normalized numerical phase velocity error as a function of direction on a Lebedev grid for biaxial dielectric. The grids cardinal axes exhibit the worst dispersion and the stability limiting axis is the one of the grid diagonals despite phase velocity depending on polarization and direction of propagation.	33
3.6	Magnetic field vector plots of the initial propagation of a pulse on a Lebedev grid with averaged sources and sampling (left) and without averaging (right). The arrows represent both x and y field components at each location because the Lebedev grid uses collocated components. The arrows in the unaveraged simulation point solely in the x and y -directions because only one of the Yee subgrids has been excited. This simulation was run in a vacuum about an electric dipole.	36
3.7	Decomposition of the Lebedev grid (center) into four offset Yee grids (surrounding). The Yee grids are displaced from one another by half the cell size in two of the x , y or z -directions.	37
3.8	The finite difference equations can be derived by approximating the line integrals along the edges of square integration loops that surround a gridpoint by the midpoint rule. On the Lebedev grid (right) there are two overlapping sets of integration loops on any given plane whereas the integration loops cover the Yee grid (left) without any overlap or gaps. Any current through a plane on the Lebedev grid passes through more than one integration loop and therefore the excitation must be distributed on the Lebedev grid.	38
3.9	Source weighting of a line of current traveling in the z -direction and passing through the center points.	42

3.10	Source weighting of an electric dipole located in the center of the grid cell. The weighting does not depend on the orientation of the dipole.	43
3.11	Source weighting about a sheet of current lying in the yz -plane and passing through the center points.	44
3.12	Setup of computational domain for simulating plane waves using only equations derived for an infinite domain. All boundaries are periodic and the excitation waveform is the same at each point so that a plane wave is launched in the $\pm z$ -directions. A short line is sampled over time and the simulation is run only long enough for a single pulse to travel through the sampled line (See Figure 3.13).	45
3.13	Plane wave pulses launched from a uniform plane source in aligned sapphire. The pulses of the two polarizations split because they have different propagation constants. A line of samples is measured over time and transformed to frequency and wavenumber to extract the propagation constant as a function of frequency for each mode (See Figure 3.14).	46
3.14	Fourier transform in time and space of sampled lines from Figure 3.13. The propagation constant as a function of frequency is found for each polarization by zooming in on the ridgeline with the chirp Z-transform at fixed frequencies. Actual results presented contain more curves because the uniform plane wave source was changed to a dipole in the center of the plane (Figure 3.15).	47
3.15	Propagation constant versus frequency for plane waves launched from an array of dipoles in three different media simulated using the Lebedev grid (solid) and the Yee grid (dashed). The data is obtained from a single simulation by Fourier transforming a line of electric field samples as described in Appendix B.1. The numerical solutions differ from the analytical solution (dotted) the least when sampled at more than 10 points per wavelength (N_λ).	49
3.16	Transverse components of the electric field for a plane wave traveling in three different media. The polarization shown in both isotropic plots is the polarization of the excitation because all polarizations have the same spatial dependence and cannot be distinguished. In anisotropic dielectrics only two linear polarizations are allowed. Their orientation depends on the direction of propagation with respect to the material's optical axis. The plots were obtained using the method in Appendix B.2.	51

4.1	Setup of computational domain for measuring the reflection coefficient of a CPML. A pulse is propagated until the reflection off the CPML at $z = L_z$ has passed the excitation plane. The CPML at $z = 0$ is three times larger than the other and is considered ideal. The reflection coefficient of each mode may be found as a function of frequency as described in Appendix B.3.	59
4.2	Reflection coefficient in dB for the two plane wave modes normally incident on a CPML (solid) and measured from a long domain (dashed). The long domain reflection coefficient represents the noise floor of the method for measuring the reflection coefficients given in Appendix B.3. The CPML has not been optimized beyond recommendations in literature but already has sufficient attenuation for the simulations in this thesis.	60
4.3	Setup of computational domain for simulating plane waves, rectangular waveguide, and parallel plate waveguide. The wave type is set by changing the boundary conditions on the x and y normal faces. A pulse is launched down the length of the domain and a line of samples is recorded. The CPML absorbing boundaries prevent unwanted reflections and allow the simulation to be run for a long time.	67
4.4	Propagation constant as a function of frequency for the modes in a rectangular waveguide simulated using the Lebedev grid (solid) and the Yee grid (dashed). The analytical solution (dotted) is known only for isotropic media. Note the TE_{mn} modes are degenerate with TM_{mn} when $m \neq 0$ and $n \neq 0$. The degeneracy is broken by the anisotropy. Differing asymptotic phase velocities cause mode lines to cross at higher frequencies. All mode lines were obtained from a single simulation run as described in Appendix B.1.	68
4.5	Transverse components of the electric field for a plane wave traveling in three different media. The basis polarization states of the uniaxial material do not align with the natural polarizations of the rectangular waveguide in the misaligned case. This conflict causes the fundamental mode to be a TE_{10} perturbed by a TE_{01} , and the second mode to be a TE_{01} perturbed by a TE_{10} . The plots were obtained using the method in Appendix B.2.	69
4.6	Grid setup for deriving the update equations at a discontinuous planar material interface. A set of grid points is assumed to lie within each material. After the integral derivation is applied, the boundary conditions are used to relate the fields on each side of the boundary. Therefore, at the interface, only the field values in material 1 need to be stored.	72

4.7	Setup of the computational domain for measuring the reflection and transmission coefficients of a dielectric interface. A pulse is launched down the length of the domain and a line of samples is recorded before and after the dielectric slab. The two line samples are transformed along z and t so that the reflection and transmission coefficients of each mode may be found as a function of frequency as described in Appendix B.3.	77
4.8	Pulse evolution over time for a plane wave normally incident on a slab of isotropic material embedded in sapphire with a permittivity tensor aligned with the grid. A line of samples is recorded in the sapphire on either side of the slab. Both lines are transformed to frequency and wavenumber to measure reflection and transmission (Figure 4.9). . . .	78
4.9	Electric field amplitude in frequency and wavenumber for plane waves normally incident on a dielectric slab. The spatial Fourier transform isolates the forwards and backwards propagating waves. The reflection and transmission coefficients are measured by zooming into each polarization at a fixed frequency using the chirp Z-transform and comparing the amplitude of the incident, reflected, and transmitted waves. . . .	79
4.10	Magnitude of the reflection (Γ) and transmission (T) coefficients for an isotropic slab ($\epsilon = 22, d = 11.8$ mm) embedded in the three media described in Section 2.3 simulated using the Lebedev grid (solid) and the Yee grid (dashed). There are two visible errors between the numerical and analytical curves (dotted): Numerical phase velocity error (downshift in frequency) and material averaging error (underestimated reflection coefficient). The Lebedev grid does not underestimate $ \Gamma $ but has greater numerical dispersion.	80
5.1	Lebedev (left) and Yee (right) grids for simulating magnetized ferrites. For both grids the magnetization vector (\vec{m}) and small signal magnetic field vector (\vec{h}) are collocated. Coupling between orthogonal components of \vec{m} and \vec{h} is easier to handle on the Lebedev grid because all field components are collocated.	85
5.2	Reduced Lebedev (right) and Yee (left) grids for simulating anisotropic dielectrics with symmetry about the z -axis. In this special case it is sufficient to represent the solution using a superposition of two Yee grids. This reduces the memory requirements of the Lebedev grid and still allows Maxwell's equations with constitutive relations to be evaluated with central differences.	87

A.1	Frequency content (top) and time domain waveform (bottom) of an excitation that is the first derivative of a Gaussian. The plots show how to determine the bandwidth and minimum time offset of the pulse. Note that the pulse has no DC content and will not deposit charge at the excitation location.	102
A.2	Frequency content (top) and time domain waveform (bottom) of a Ricker wavelet. The plots show how to determine the bandwidth and minimum time offset of the pulse. Note that the pulse has no DC content and will not deposit charge at the excitation location.	103
A.3	Frequency content (top) and time domain waveform (bottom) of a cosine modulated Gaussian used for narrow band excitations. The plots show how to determine the bandwidth and minimum time offset of the pulse. The DC content of the waveform goes to zero when the center frequency is higher than the pulse bandwidth.	104
A.4	Spatial distribution (top) and excitation amplitude of different modes (bottom) for a sine windowed Gaussian (solid) and an infinitesimal dipole (dotted). The parameters of the sine windowed Gaussian are $x_0/L_x = 0.2$ and $\sigma/L_x = 0.3$	105
A.5	Magnitude of electric field components versus wavenumber at a fixed frequency for a rectangular waveguide. The dipole excitation (top) contains artifacts which are significantly reduced when a spatially distributed excitation (bottom) is used. The peaks are labeled with the FDTD-analytic solution. Note that the TH modes are degenerate with TE modes.	106
C.1	Structure of the Lebedev grid. The electric and magnetic fields each form an FCC grid but together they form an easily indexed Cartesian grid. All field components are collocated. The update stencil only extends to the immediately adjacent gridpoints of opposite field type.	114

List of Abbreviations and Symbols

Abbreviations

FDTD	Finite Difference Time Domain
PML	Perfectly Matched Layer
PEC	Perfect Electric Conductor
FCC	Face Centered Cubic
DFT	Discrete Fourier Transform

Symbols

Δt	Timestep
n	Timestep index
$\Delta x, \Delta y, \Delta z$	Cell sizes in each direction
i, j, k	Gridpoint indices in each direction
\vec{E}	Electric field
\vec{H}	Magnetic field
\vec{D}	Electric displacement
\vec{B}	Magnetic flux density
\vec{J}	Impressed electric current
\vec{M}	Impressed magnetic current
$\bar{\epsilon}$	Permittivity tensor
$\bar{\mu}$	Permeability tensor
$\bar{\sigma}$	Conductivity tensor
$\bar{\sigma}^*$	Magnetic loss tensor
δ_{ij}	Kronecker-Delta function
ε_{ijk}	Levi-Civita function
λ	Wavelength
\vec{k}	Propagation constant
f	Frequency

Chapter 1

Introduction

1.1 Motivation

The finite-difference time-domain (FDTD) method has become a valuable tool for solving large electromagnetic problems and mature simulation tools are commercially available. However, the problems being considered by scientists and engineers continue to push the limits of what these simulation tools can handle. In particular, novel materials are being incorporated into circuit design that require a more general mathematical formulation than traditional materials. The standard discretization of the FDTD equations cannot easily handle such complicated materials. For this reason, the primary goal of this thesis is to explicitly show how to develop finite-difference algorithms on non-standard grids beginning with the propagation of plane waves and through to analyzing transmission lines that contain metal, material, and absorbing boundary conditions. This is done by developing an algorithm for simulating anisotropic media on a so-called Lebedev grid. However, the derivations and problem approaches can be generalized to other grids as well.

Commercial FDTD implementations are based on the Yee grid [1]. This discretization has proven to be robust because it is the most sparse grid using only central differences to approximate the derivatives in Maxwell's curl equations. How-

ever, this is only true for isotropic media. When general anisotropy is introduced the tensor constitutive relations couple in spatial derivatives that do not normally need to be evaluated. They cannot be evaluated by central differences on the Yee grid because the electric and magnetic field components are not stored in the correct locations. Current methods for handling anisotropy require interpolations that are computationally expensive, break the simple nature of the updating algorithm, reduce the order of accuracy of the simulation, and/or cause instabilities at material and metal interfaces [2, 3, 4].

The proposed solution to this problem is to change the grid discretization to the most sparse grid that uses only central differences to approximate any spatial or temporal derivatives in the describing differential equations. For anisotropic media, this grid is known as a Lebedev grid [5] and is used to solve the elastic wave equations for seismic imaging [6, 7]. The grid was originally proposed by Lebedev in 1964 [8, 9], but did not find use until recent years because the computational cost of full wave simulations was too high. It has been used in electromagnetic finite-difference frequency-domain simulations by Davydycheva [5] and was independently proposed for FDTD by Garcia [2]. Although the basics of the grid can be found in literature, there is no work describing how to use the grid for simulating microwave transmission lines and antennas, nor is there a fully anisotropic dispersion relation with proof of stability and conservation of charge. The only dispersion relations in literature are for isotropic media simulated on a Lebedev grid [2] and for elastodynamics [7].

1.2 Thesis Goals

In summary, the goals of this thesis are as follows:

- To show how to derive a grid from the central difference based discretization of a set of differential equations using anisotropic media as an example.
- To prove that the algorithm is stable and to investigate the expected computational errors.
- To implement planar PEC, CPML, and material interfaces on the grid while maintaining an algorithm that is parallelizable and simple to program.

1.3 Thesis Outline

The work provided in this thesis complements the work on the Lebedev grid that is already found in literature and highlights the reasons that the Lebedev grid is the natural extension to the Yee grid. The derivations are intertwined with numerical results that show the relevance and prove the functionality of each step towards producing a comprehensive simulation tool. After a brief introduction to anisotropy and FDTD in Chapter 2 the thesis is organized as follows:

- Chapter 3 describes how to handle a grid of infinite extent containing a single anisotropic material. This basic case is used to derive the formulas that hold in the bulk of any material, prove the stability of the recursive algorithm, and examine how sources need to be implemented on the grid.
- Chapter 4 adds some necessary components of a practical simulation tool for analyzing microwave transmission lines and antennas. In particular, methods for handling multiple materials and perfect conductors are derived. The implementation of a convolutional perfectly matched layer (CPML) absorbing boundary condition is also given.
- Chapter 5 shows how to decide on the form of a grid given a set of differen-

tial equations. In particular, an extension of the Lebedev grid is given that allows the auxiliary equations describing the interactions within a ferrite to be approximated alongside Maxwell's equations and a special grid that requires less memory than the Lebedev grid is provided for the special case of uniaxial dielectrics.

The conclusion in Chapter 6 provides a summary of the accomplishments of this thesis, states the caveats that have been learned, and suggests some avenues for continued development. The appendices may also be of interest, even to experts in the FDTD method. Appendix A demonstrates the importance of setting up the frequency content and spatial distribution of the excitation. Appendix B provides a method of extracting all of the propagation constants of the allowed modes in a transmission line as a function of frequency using a single simulation run. This method is extended to measure reflection and transmission coefficients. Appendix B also covers how to plot the cross-section of each mode without storing an overwhelming amount of data for post-processing. Finally, Appendix C describes an object-oriented approach to setting up a computational domain. It also describes how to take advantage of the symmetry in the governing equations to avoid error prone copy-and-paste style coding.

Chapter 2

Background

The FDTD method is a computational scheme that discretizes the fields onto a regular mesh and uses a recursive equation for each gridpoint to march the solution forward from the initial conditions. The recursive update at each gridpoint depends only on a small number of neighboring points. Furthermore, the equation takes the same form for every point in the simulation and even has the same coefficients for every point in the same material. This similarity between formulas allows extremely fast and parallel hardware implementations of the algorithm. It also allows the programmer to use modern object oriented concepts such as polymorphism and inheritance to simplify the software that sets up the computational domain. Maintaining these benefits while incorporating complex electromagnetic materials is an underlying theme in this work.

Anisotropic media refers to materials that have different properties in different directions. A familiar example of a mechanically anisotropic material is wood, which has different properties with and against the grain. In electromagnetics, naturally occurring materials such as crystals, bodily tissues and engineered materials, such as circuit board substrates, all exhibit anisotropy [10, 11]. There are several subclasses of anisotropic materials including anisotropic dielectrics, anisotropic conductors, and gyrotropic media [12]. Synthetic sapphire is an example of an anisotropic dielectric

that is used in high frequency integrated circuits because it has low electromagnetic loss and good thermal properties [10]. Anisotropic conductors include muscles and arrays of conducting plates or wires that are engineered to act as linear polarizers. Finally, examples of gyrotropic media are ferrites and magnetized plasma.

All of these materials have the same mathematical representation. That is, at least one of the permittivity ($\bar{\epsilon}$), permeability ($\bar{\mu}$), conductivity ($\bar{\sigma}$), or magnetic loss ($\bar{\sigma}^*$) must be represented by a tensor instead of a scalar. The algorithms developed for the Lebedev grid assume that all of the constitutive relations are described by tensors so that any of the materials mentioned previously can be included. Note that gyrotropic materials are typically strongly dispersive but can be approximated by a frequency independent tensor over a narrow range of frequencies. Better models of gyrotropic materials take into account an auxiliary set of differential equations that describe their behavior. An example of how to handle the additional equations is given in Chapter 5.

2.1 Governing Equations

The governing equations in FDTD are Maxwell's curl equations and the constitutive relations. The first equation states that a magnetic current density (\vec{M}) or a change in magnetic flux density (\vec{B}) will induce a solenoidal electric field (\vec{E}). Although the magnetic current density is a theoretical quantity, it can be useful in simulating equivalent circuits:

$$\nabla \times \vec{E} = -\frac{\partial \vec{B}}{\partial t} - \vec{M}_{\text{total}} \quad . \quad (2.1)$$

This equation is commonly known as Faraday's law. The second equation, known as Ampere's law, states that a current density (\vec{J}) or a change in the electric displacement (\vec{D}) will induce a solenoidal magnetic field (\vec{H}):

$$\nabla \times \vec{H} = \frac{\partial \vec{D}}{\partial t} + \vec{J}_{\text{total}} \quad . \quad (2.2)$$

The constitutive relations are much less clear cut and will change with the assumed material model. In general, \vec{D} and \vec{B} are found from \vec{E} and \vec{H} using material parameters. These parameters may depend on frequency (dispersive), direction (anisotropic), position (inhomogeneous), the magnitude of the fields (non-linear), or external variables such as temperature and pressure (multiphysics problems). This work considers anisotropic materials, in which case

$$\begin{aligned} \vec{D} &= \bar{\epsilon} \vec{E} \quad , \\ \vec{B} &= \bar{\mu} \vec{H} \quad , \end{aligned} \quad (2.3)$$

where $\bar{\epsilon}$ and $\bar{\mu}$ are tensors.

It is also common for materials to have a finite conductivity ($\bar{\sigma}$). That is, a current will flow when there is an applied electric field. Like the permittivity and permeability, the conductivity depends on many parameters in different materials. When only anisotropy is considered, Ohm's law becomes

$$\vec{J}_{\text{total}} = \vec{J} + \bar{\sigma} \vec{E} \quad . \quad (2.4)$$

Finite conductivities at high frequencies will result in losses within the material. The losses can also be proportional to the propagating magnetic field. Magnetic losses

$(\bar{\sigma}^*)$ are expressed in an analogous manner by assuming

$$\vec{M}_{\text{total}} = \vec{M} + \bar{\sigma}^* \vec{H} \quad . \quad (2.5)$$

Here \vec{J} and \vec{M} are known as the impressed currents and are used as sources.

There are two more governing equations that need to be taken into account. Gauss' Law states that the electric displacement originates from and terminates on electric charge (ρ_e):

$$\nabla \cdot \vec{D} = \rho_e \quad . \quad (2.6)$$

The analogous law for magnetic fields is

$$\nabla \cdot \vec{B} = \rho_m \quad . \quad (2.7)$$

In reality, there is no such thing as magnetic charge, $\rho_m = 0$, but like the magnetic current, including magnetic charge can be useful when using equivalence theorems.

The divergence equations are rarely used for deriving FDTD update equations, but they are still extremely important. The standard Yee implementation [13], extensions such as the Lobatto cell for inhomogeneous material by Chilton [14], and the algorithm presented in this thesis each satisfy both divergence equations. Algorithms that do not satisfy these laws can still be useful, but will contain artifacts from unphysical charge buildup at the gridpoints.

The equations used to derive the FDTD update equations are a combination of Maxwell's curl equations, the anisotropic constitutive relations, and the anisotropic

version of Ohm's law such that the fields are written in terms of \vec{E} , \vec{H} , \vec{J} , and \vec{M} :

$$\nabla \times \vec{H} = \bar{\epsilon} \frac{\partial \vec{E}}{\partial t} + \bar{\sigma} \vec{E} + \vec{J} \quad , \quad (2.8)$$

$$\nabla \times \vec{E} = -\bar{\mu} \frac{\partial \vec{H}}{\partial t} - \bar{\sigma}^* \vec{H} - \vec{M} \quad . \quad (2.9)$$

The symmetry of these two equations is essential to both derivations and the programming structure adopted. Note that the equations can be interchanged by the transformation

$$\begin{aligned} \vec{E} &\rightarrow \vec{H} \quad , \\ \bar{\epsilon} &\rightarrow -\bar{\mu} \quad , \\ \bar{\sigma} &\rightarrow -\bar{\sigma}^* \quad , \\ \vec{J} &\rightarrow -\vec{M} \quad . \end{aligned}$$

This symmetry is known as duality.

Of equal importance is the integral form of Maxwell's curl equations because they provide a more physically intuitive derivation of the FDTD scheme that is easier to generalize to sources and material boundaries. This form is found by integrating the flux through a surface and using Stoke's theorem to reduce the curl term from a surface integral to a contour integral:

$$\oint \vec{H} \cdot d\vec{L} = \iint \left(\bar{\epsilon} \frac{\partial \vec{E}}{\partial t} + \bar{\sigma} \vec{E} + \vec{J} \right) \cdot d\vec{S} \quad , \quad (2.10)$$

$$\oint \vec{E} \cdot d\vec{L} = \iint \left(-\bar{\mu} \frac{\partial \vec{H}}{\partial t} - \bar{\sigma}^* \vec{H} - \vec{M} \right) \cdot d\vec{S} \quad . \quad (2.11)$$

These equations are the basis for the more complicated derivations seen in Chapter 4.

2.2 Electromagnetic Effects of Anisotropic Dielectrics

In most cases, simulations are conducted because the problem is too complicated for analytical techniques. However, there are many things that can go wrong in both derivations and implementations of numerical methods. Therefore, understanding the basic consequences of anisotropy is necessary for interpreting and verifying simulation results. Most practical transmission lines do not have analytical solutions when they are filled with anisotropic materials. That said, the transmission line solutions are well known for the isotropic case and the effects of anisotropy can be viewed as a perturbation on those solutions.

With this in mind, the goal of this section is to develop an intuition for how anisotropy will change familiar isotropic solutions. In the end, it all comes back to plane waves. Understanding how anisotropy affects plane waves is the first step towards predicting its effects on wave guiding structures. A textbook with a comprehensive section on propagation in anisotropic media is written by Someda [12]. Collin [15, 16], and Chew [17] also provide useful insight, but do not provide as broad of coverage.

The main consequence of anisotropy is a splitting of degenerate modes. For plane waves in isotropic media, all waves propagate with the same phase velocity in any direction and with any polarization. Anisotropy breaks this degeneracy. In anisotropic

media, there are two eigensolutions to the plane wave equation that each have different propagation constants. The polarizations of the two solutions are orthogonal to one another but not necessarily orthogonal to the direction of propagation. They must be used as a basis to express any other polarization state. Since the propagation constants of these two basis states differ, the polarization state of any linear combination will necessarily be a function of position. Faraday rotation is a particular instance of this phenomenon. In gyrotropic media, there is an axis along which the two basis states are circularly polarized waves. When a linear polarization passes through the gyrotropic medium the polarization state must be expressed as two circularly polarized waves with different wavelengths. The total polarization remains linear, but its orientation rotates in space.

The basis polarization states and their propagation constants (\vec{k}) can be found in a relatively straightforward manner. The derivation of how to find them is important because it is similar to the derivation of the FDTD dispersion relation given later. The technique is a modification to the plane wave velocity derivation found in [12] for lossless dielectrics. The difference is that a more general eigenvalue equation is developed because it is not assumed that the material is a lossless dielectric, and that the eigenvalue equation is solved numerically once the material parameters have been chosen.

Consider a plane wave propagating with angular frequency ω :

$$\vec{E} = \vec{E}_0 \exp(j(\omega t - \vec{k} \cdot \vec{r})) \quad , \quad (2.12)$$

$$\vec{H} = \vec{H}_0 \exp(j(\omega t - \vec{k} \cdot \vec{r})) \quad . \quad (2.13)$$

In this case the curl reduces to a cross product:

$$\nabla \times \vec{E} = -j\vec{k} \times \vec{E} \quad , \quad (2.14)$$

and the time derivative reduces to a scalar multiplication:

$$\frac{\partial \vec{E}}{\partial t} = j\omega \vec{E} \quad . \quad (2.15)$$

In the absence of sources, Maxwell's equations for plane waves become

$$\nabla \times \vec{E} = -\bar{\mu} \frac{\partial \vec{H}}{\partial t} \quad \Longrightarrow \quad -j\vec{k} \times \vec{E} = -j\omega \bar{\mu} \vec{H} \quad , \quad (2.16)$$

$$\nabla \times \vec{H} = \bar{\epsilon} \frac{\partial \vec{E}}{\partial t} \quad \Longrightarrow \quad -j\vec{k} \times \vec{H} = j\omega \bar{\epsilon} \vec{E} \quad . \quad (2.17)$$

It is convenient to split the propagation constant into its magnitude (k) and direction (\hat{r}). Note that the cross product between two vectors can be equivalently represented by a matrix-vector multiply:

$$\hat{r} \times \vec{E} = \begin{pmatrix} 0 & -r_z & r_y \\ r_z & 0 & -r_x \\ -r_y & r_x & 0 \end{pmatrix} \vec{E} = \bar{\bar{r}} \vec{E} \quad . \quad (2.18)$$

Equations (2.16) and (2.17) can be combined to give

$$-\bar{\epsilon}_r^{-1} \bar{\bar{r}} \bar{\mu}_r^{-1} \bar{\bar{r}} \vec{E} = \frac{\omega^2 \epsilon_0 \mu_0}{k^2} \vec{E} = \frac{c^2}{c_0^2} \vec{E} \quad , \quad (2.19)$$

where $c = \omega/k$ is the phase velocity, $c_0 = 1/\sqrt{\epsilon_0 \mu_0}$ is the speed of light in a vacuum, $\bar{\epsilon}_r$

is the relative permittivity tensor, and $\bar{\bar{\mu}}_r$ is the relative permeability tensor. Equation (2.19) is an eigenvalue equation with eigenvalue c^2/c_0^2 . Therefore, the eigenvalues of the matrix $-\bar{\bar{\epsilon}}_r^{-1}\bar{\bar{r}}\bar{\bar{\mu}}_r^{-1}\bar{\bar{r}}$ are related to the phase velocity of each allowed polarization and depend on the direction of propagation. It is convenient to write this in terms of the refractive index (n):

$$n = \frac{c_0}{c} = \frac{1}{\sqrt{\text{eig}(-\bar{\bar{\epsilon}}_r^{-1}\bar{\bar{r}}\bar{\bar{\mu}}_r^{-1}\bar{\bar{r}})}} \quad . \quad (2.20)$$

This equation has been solved as a function of direction for several example materials and is plotted in Figures 2.1 to 2.4.

The eigenvectors corresponding to the eigenvalues can be normalized to give the basis electric field polarization states in the media. The magnetic field polarization state is then given by

$$\vec{H} = \frac{c_0}{n} \bar{\bar{\mu}}^{-1} \bar{\bar{r}} \vec{E} \quad . \quad (2.21)$$

There are constraints on the material parameter tensors that must be satisfied in order for this method to work. However, the exact form of these constraints is unclear. Suffice it to say that if the input tensors represent a physical medium that supports plane waves then there will be two non-zero eigenvalues and one zero eigenvalue. This can be proven for a lossless biaxial electric or magnetic material but the formula is more general than this. In fact, for gyrotropic and conductive media the natural extension is $\bar{\bar{\epsilon}}_r \rightarrow \bar{\bar{\epsilon}}_r + j\bar{\bar{\sigma}}/\omega\epsilon_0$ and $\bar{\bar{\mu}}_r \rightarrow \bar{\bar{\mu}}_r + j\bar{\bar{\sigma}}^*/\omega\mu_0$.




Figure 2.1: Plane wave velocities as a function of direction for an isotropic dielectric. The phase velocity is given by a vector between the origin and a point on the surface for each polarization. In isotropic media, any polarization has the same phase velocity that is independent of direction.




Figure 2.2: Plane wave velocities of the two basis polarizations as a function of direction for a uniaxial dielectric. The phase velocity is given by a vector between the origin and a point on the surface for each polarization. In a uniaxial dielectric, there are two basis linear polarization states in any direction of propagation. The phase velocity of one state is independent of direction (ordinary wave) whereas the other depends on direction (extraordinary wave). There is a single axis (optical axis) where the phase velocities for both states are the same. Note that the anisotropy is exaggerated for the sake of illustration.




Figure 2.3: Plane wave velocities of the two basis polarizations as a function of direction for a biaxial dielectric. The phase velocity is given by a vector between the origin and a point on the surface for each polarization. In a biaxial dielectric, there are two basis linear polarization states in any direction of propagation and their phase velocities depend on direction. There are two axes (optical axes) where the phase velocities for both states are the same. Note that the anisotropy is exaggerated for the sake of illustration.




Figure 2.4: Plane wave velocities of the two basis polarizations as a function of direction for a gyrotropic dielectric. The phase velocity is given by a vector between the origin and a point on the surface for each polarization. In a gyrotropic dielectric, there are two basis elliptical polarization states in any direction of propagation. The polarization states are of different handedness for opposite directions of propagation. This leads to non-reciprocal behavior.

2.3 Sapphire Permittivity Tensor

When developing a simulation tool it is essential to compare the results of numerical simulations with known analytical solutions. Unfortunately, analytical solutions are relatively rare for general anisotropy. That said, anisotropic dielectrics are a simple class of anisotropic materials that are both well understood, and exhibit many of the phenomena associated with general anisotropy. Furthermore, anisotropic dielectrics are commonly encountered in practice. Therefore, the numerical simulations in this thesis will be focused on anisotropic dielectrics.

In particular, numerical examples throughout the thesis will consider wave propagation in sapphire. Sapphire is a crystal with the chemical formula Al_2O_3 that has a low loss tangent at microwave frequencies and good heat conducting properties. Synthetic sapphire is commercially available and can be manufactured with a specified crystal orientation [10]. The permittivity tensor for sapphire in its principal axis frame is [18]

$$\bar{\bar{\epsilon}}_{\text{Aligned}} = \begin{pmatrix} 11.54 & 0 & 0 \\ 0 & 9.34 & 0 \\ 0 & 0 & 9.34 \end{pmatrix} \epsilon_0 \quad . \quad (2.22)$$

In this case the optical axis is in the x -direction. For arbitrary orientations of the optical axis the permittivity tensor depends on the rotation matrix ($\bar{\bar{R}}$) of the transformation as

$$\bar{\bar{\epsilon}}_{\text{Rotated}} = \bar{\bar{R}} \bar{\bar{\epsilon}} \bar{\bar{R}}^{-1} \quad . \quad (2.23)$$

The numerical examples use a permittivity that has been rotated by 32° about the

z -axis and 11° about the y -axis:

$$\bar{\bar{\epsilon}}_{\text{Misaligned}} = \begin{pmatrix} 10.86 & 0.97 & -0.30 \\ 0.97 & 9.96 & -0.19 \\ -0.30 & -0.19 & 9.40 \end{pmatrix} \epsilon_0 \quad . \quad (2.24)$$

The numerical simulations also use an isotropic approximation to the sapphire permittivity tensor that can be compared with analytical solutions. The scalar permittivity is taken to be the mean of the diagonal elements in the principal axis reference frame

$$\bar{\bar{\epsilon}}_{\text{Isotropic}} = \begin{pmatrix} 10.07 & 0 & 0 \\ 0 & 10.07 & 0 \\ 0 & 0 & 10.07 \end{pmatrix} \epsilon_0 \quad . \quad (2.25)$$

2.4 FDTD Techniques for Handling Anisotropy

Regular FDTD based on the Yee grid can only handle anisotropy when the material's principal axis frame aligns with the FDTD coordinate system [13]. Essentially, the diagonal elements of the permittivity tensor scale the equations in the same way as an unequal aspect ratio, $\Delta x \neq \Delta y \neq \Delta z$. However, if the FDTD coordinate frame does not align with the material principal axis frame then the off-diagonal elements of the tensors become non-zero. Full tensors couple in extra spatial derivatives from the curl equations that do not normally need to be evaluated. Therefore, the regular FDTD method must be extended [19].

The oldest approach is to derive a different stencil for these derivatives on the regular Yee grid [20, 2]. The derivation in [20] begins by discretizing the derivatives with

central differences. The terms that fall between gridpoints are then evaluated using a four point interpolation. It follows that each non-standard derivative is evaluated with an 8 point stencil instead of the 2 point stencil of a central difference. These large stencils significantly reduce the speed of the algorithm because accessing memory is the limiting factor in most FDTD simulations. Furthermore, the form of these interpolations becomes unclear when implementing boundary conditions. Consequently, interpolation methods are notorious for going late-time unstable at boundaries [3].

A second approach is to store \vec{D} collocated with \vec{E} , and \vec{B} collocated with \vec{H} [3, 21]. These schemes double the required memory storage but are capable of handling PEC boundary conditions [3, 4] and material interfaces [21, 22, 23]. The downside is that these methods are both complicated and use a large stencil. Complicated methods are more susceptible to programming errors and also have a narrower scope of application. Additionally, large stencils increase the simulation time because compute time is limited by memory bottlenecks.

The literature method compared to in the numerical examples is based on collocating \vec{D} and \vec{E} . Specifically, the interpolation scheme given by Werner and Cary is used to interpolate E from D [21]. Werner and Cary proposed a two step interpolation for inhomogeneous materials that achieves second order error for slowly varying dielectrics and first order error for discontinuous dielectrics. A material averaging scheme proposed by Oskooi was adopted to improve the error at dielectric discontinuities [22]. PEC was implemented using the extrapolation scheme by Zhao [3] based on the scheme proposed for ferrites by Okoniewski [24].

The final approach is the one pursued in this thesis. That is, store field components in such a way that all derivatives are evaluated using central differences. At first

glance, this method has an increased cost in memory, but does not require large stencil interpolations [2]. It has also been shown that in certain cases the increase in accuracy relieves the memory cost because a more sparse grid can achieve the same level of error [5]. Furthermore, the Lebedev algorithm is based on a rigorous integral derivation as opposed to interpolations of differential stencils.

Chapter 3

Lebedev Grid in an Unbounded Domain

The first step in developing a simulation tool is to solve Maxwell's equations in a single material without any boundaries. Although this may seem like an extremely limiting case, the right domain setup allows the method to be tested against the analytical solution for plane waves and is also easy to implement. This is important because bugs in the simulation code often lead to instabilities and it becomes unclear whether the instability is a theoretical problem or an implementation problem. Furthermore, this setup can be used to test the stability limit without worrying whether instabilities are caused by particular features within the domain. Once the basic stability limit is well understood, additional features can be added as required.

Consequently, the chapter is organized as follows. First, the notation is described and the update equations are derived using central differences. Then, the grid structure is derived from the resulting discretized equations. After that, the stability criterion is derived and tested. At this point some nuances will arise that need to be dealt with. Namely, the Lebedev grid can be decomposed into four Yee grids that must be appropriately coupled together. An alternative derivation based on the integral form of Maxwell's equations shows how the coupling can be implemented. After

those details are dealt with, the simulation is setup with periodic boundaries (which do not change the form of the update equations) and used to simulate plane waves. Finally, the chapter will conclude by proving that the method satisfies the divergence equations because the Lebedev grid has a consistent calculus.

3.1 Notation

The update equations that follow contain a large number of terms. However, these terms also contain a high degree of symmetry. This symmetry is exploited by writing the equations using sums, Kronecker delta functions (δ_{ij}), and Levi-Civita functions (ε_{ijk}). The notation used in this thesis is slightly different than the notation found in literature. The notation is adopted because it is both concise and flows naturally with the programming structure described in Appendix C. If the grid index increases twice as fast as the cell size then the Lebedev grid can be stored in a single $3 \times 2N_x \times 2N_y \times 2N_z$ array. The three components of the electric fields are stored wherever $i + j + k$ is even and the magnetic fields where their sum is odd. Unlike the Yee grid, all components have the same relation between index and spatial location ($x = i\Delta x/2$). This simplifies the relationship between neighboring fields and the notation is intended to intuitively match this storage structure.

The superscript on a variable denotes the index of the timestep. For the n^{th} timestep, \vec{E}^{+1} is evaluated at time $t = (n + 1)\Delta t$. Square brackets following a vector denote its relative location to the vector being updated. For example, if $\vec{E}(i\Delta x/2, j\Delta y/2, k\Delta z/2)$ is being updated it is denoted $\vec{E}[0, 0, 0]$ and the magnetic field one half cell over in the negative y -direction is denoted $\vec{H}[0, -1, 0]$. Two special shorthands are $\vec{E}[0] = \vec{E}[0, 0, 0]$ and $\vec{E}[\delta_w] = \vec{E}[\delta_{wx}, \delta_{wy}, \delta_{wz}]$. Subscripts on a tensor

such as ϵ_{uv} denote the element in the u^{th} row and v^{th} column of $\bar{\epsilon}$. Square brackets following a tensor select a single tensor from a set of tensors. For example, $\bar{c}[+, x]$ is the tensor that premultiplies the magnetic field vector one grid cell in the positive x -direction from the electric field being updated. The sums range through the elements in the sequence below summation symbol. For example, if the subscript is $v = xyz$ then the terms following the sum are added together with $v = x$, $v = y$, and $v = z$.

3.2 Central Difference Based Derivation of the Update Equations

The FDTD method is based on the differential form of Maxwell's curl equations, (2.8) and (2.9), repeated here for convenience:

$$\nabla \times \vec{H} = \bar{\epsilon} \frac{\partial \vec{E}}{\partial t} + \bar{\sigma} \vec{E} + \vec{J} \quad , \quad (3.1)$$

$$\nabla \times \vec{E} = -\bar{\mu} \frac{\partial \vec{H}}{\partial t} - \bar{\sigma}^* \vec{H} - \vec{M} \quad . \quad (3.2)$$

The derivation for the electric field updates follows the first steps of the derivation in [20]. Consider the u^{th} row of (3.1) in the absence of sources:

$$\left(\nabla \times \vec{H} \right)_u^{+\frac{1}{2}} = \sum_{v=xyz} \left(\epsilon_{uv} \frac{\partial E_v^{+\frac{1}{2}}}{\partial t} + \sigma_{uv} E_v^{+\frac{1}{2}} \right) \quad . \quad (3.3)$$

A staggered in time scheme is ideal because it halves the required memory. Therefore, \vec{E} is only available at integer timesteps. The time derivative is approximated by a

central difference:

$$\frac{\partial E_v^{+\frac{1}{2}}}{\partial t} \approx \frac{E_v^{+1} - E_v^0}{\Delta t} \quad , \quad (3.4)$$

and the value between the integer timesteps is replaced by a central average:

$$E_v^{+\frac{1}{2}} \approx \frac{E_v^{+1} + E_v^0}{2} \quad . \quad (3.5)$$

This gives

$$\left(\nabla \times \vec{H} \right)_u^{+\frac{1}{2}} = \sum_{v=xyz} \left(\epsilon_{uv} \frac{E_v^{+1} - E_v^0}{\Delta t} + \sigma_{uv} \frac{E_v^{+1} + E_v^0}{2} \right) \quad . \quad (3.6)$$

The curl, which can be written concisely using the Levi-Civita symbol (ϵ_{uvw}):

$$\epsilon_{uvw} = \begin{cases} 1 & \text{if } (u, v, w) = (x, y, z) \text{ or } (y, z, x) \text{ or } (z, x, y) \\ -1 & \text{if } (u, v, w) = (x, z, y) \text{ or } (z, y, x) \text{ or } (y, x, z) \\ 0 & \text{else} \end{cases} \quad , \quad (3.7)$$

$$\left(\nabla \times \vec{H} \right)_u = \sum_{v=xyz} \sum_{w=xyz} \epsilon_{uvw} \frac{\partial H_v}{\partial w} \quad , \quad (3.8)$$

is also discretized using central differences:

$$\frac{\partial H_v}{\partial w} [0] \approx \frac{H_v[\delta_w] - H_v[-\delta_w]}{\Delta w} = \sum_{p=\pm 1} \frac{p H_v[p\delta_w]}{\Delta w} \quad . \quad (3.9)$$

This gives

$$\left(\nabla \times \vec{H} \right)_u^{+\frac{1}{2}} = \sum_{v=xyz} \sum_{w=xyz} \sum_{p=\pm 1} \epsilon_{uvw} \frac{p H_v^{+\frac{1}{2}}[p\delta_w]}{\Delta w} \quad . \quad (3.10)$$

Equations (3.6) and (3.10) are then combined:

$$\sum_{v=xyz} \left(\epsilon_{uv} \frac{E_v^{+1}[0] - E_v^0[0]}{\Delta t} + \sigma_{uv} \frac{E_v^{+1}[0] + E_v^0[0]}{2} \right) = \sum_{v=xyz} \sum_{w=xyz} \sum_{p=\pm 1} \epsilon_{uvw} \frac{p H_v^{+\frac{1}{2}}[p \delta_w]}{\Delta w} . \quad (3.11)$$

This equation holds for any u . It can be rearranged into a matrix equation where u is the row index and v is the column index:

$$\bar{a} \vec{E}^{+1}[0] = \bar{b} \vec{E}^0[0] + \sum_{p=\pm 1} \sum_{w=xyz} \bar{c}[p, w] \vec{H}^{+\frac{1}{2}}[p \delta_w] . \quad (3.12)$$

Where

$$a_{uv} = \frac{\epsilon_{uv}}{\Delta t} + \frac{\sigma_{uv}}{2} , \quad (3.13)$$

$$b_{uv} = \frac{\epsilon_{uv}}{\Delta t} - \frac{\sigma_{uv}}{2} , \quad (3.14)$$

$$c_{uv}[p, w] = \epsilon_{uvw} \frac{p}{\Delta w} . \quad (3.15)$$

An analogous equation of the magnetic field update can be derived or found using duality:

$$\bar{\alpha} \vec{H}^{+\frac{1}{2}}[0] = \bar{\beta} \vec{H}^{-\frac{1}{2}}[0] + \sum_{p=\pm 1} \sum_{w=xyz} \bar{\gamma}[p, w] \vec{E}^0[p \delta_w] . \quad (3.16)$$

Where

$$\alpha_{uv} = \frac{-\mu_{uv}}{\Delta t} - \frac{\sigma_{uv}^*}{2} , \quad (3.17)$$

$$\beta_{uv} = \frac{-\mu_{uv}}{\Delta t} + \frac{\sigma_{uv}^*}{2} , \quad (3.18)$$

$$\gamma_{uv}[p, w] = \epsilon_{uvw} \frac{p}{\Delta w} . \quad (3.19)$$

Note that \bar{a} and $\bar{\alpha}$ are both tensors that depend on material properties and are invertible for any physical material. It is easiest to program in the values in this form, invert \bar{a} and $\bar{\alpha}$ numerically and multiply them through the right-hand-side. The inversion can be done analytically, but it is unnecessary because the overhead cost of a handful of 3 by 3 matrix inversions is negligible.

3.3 Grid Structure

At this point the observant reader will notice that the grid must have collocated field components and that the electric and magnetic fields are staggered in time and space. For example, in (3.12), the electric field at the current timestep depends on the electric field at the same location at the previous timestep and magnetic fields one half cell away in each direction at the previous half timestep. At each gridpoint all of the components must be stored because in general $\bar{a}^{-1}\bar{b}$ and $\bar{\alpha}^{-1}\bar{\beta}$ are full matrices.

The resulting grid, known as a Lebedev grid, is plotted in Figure 3.1 with an overlay of the adopted notation. The grid can be viewed from several different perspectives. First of all, the Lebedev grid is a regular Cartesian lattice with collocated field components and alternating points representing electric and magnetic fields. Alternatively it can be viewed as two shifted face-centered cubic (FCC) grids. There is one FCC grid for the electric field and one for the magnetic field. To construct this grid, number the gridpoints in the x , y , and z -directions with indices i , j , and k . If $i + j + k$ is even, all the components of \vec{E} are stored, and if $i + j + k$ is odd, all the components of \vec{H} are stored [5].

Figure 3.2 contrasts the Lebedev grid with the Yee grid. It is immediately apparent that the Lebedev grid requires more storage for the same cell size. However, modern




Figure 3.1: Indexing notation for the Lebedev grid. These fields are indexed with respect to the electric field located at $(i, j, k) = (2, 1, 3)$. Together the electric and magnetic fields form a Cartesian grid which allows the fields to be stored together in a regular 3D storage format.

Yee grid based anisotropic codes use larger stencils and therefore have a higher compute cost in every update. Furthermore, the Lebedev grid only uses central differences whereas Yee grid based approaches interpolate. There are likely situations where the Lebedev grid can achieve the same level of accuracy with a larger cell size. Finally, there is no formula for the maximum stable timestep in Yee grid based approaches for anisotropic media. Therefore, a low estimate is used and a greater number of timesteps is required to propagate the solution for the same amount of time.

3.4 Dispersion and Stability Analysis

A numerical algorithm is considered stable when a bounded source function excites bounded fields within the domain. Explicit FDTD methods are typically stable only when an inequality relating the timestep and the cell spacing is satisfied. Knowing




Figure 3.2: Traditional Yee grid (left) and proposed Lebedev grid (right). The Lebedev grid uses collocated field components and staggered field types so that central derivatives can be used to compute the spatial derivative of any field component along any axis. The Yee grid uses staggered field components and staggered field types because it is specialized to compute only the spatial derivatives required for Maxwell's equations in isotropic materials.

the limit of this inequality allows the user to run the simulation the furthest in time with the least amount of resources. The dispersion relation is tied to the stability analysis. It predicts the error in the propagation constants of numerical plane waves as a function of direction of travel. Knowledge of this error helps explain other observed discrepancies between FDTD simulations and analytical or measured results.

The dispersion relation derivation follows the method of obtaining analytic plane wave velocities given in Section 2.2. Consider a plane wave propagating with angular frequency ω and propagation constant \vec{k} :

$$\vec{E} = \vec{E}_0 \exp(j(\omega t - \vec{k} \cdot \vec{r})) \quad , \quad (3.20)$$

$$\vec{H} = \vec{H}_0 \exp(j(\omega t - \vec{k} \cdot \vec{r})) \quad . \quad (3.21)$$

On the grid: $t = n\Delta t$, $x = i\Delta x/2$, $y = j\Delta y/2$, and $z = k\Delta z/2$. The matrix form of (3.11) describes the relation between \vec{E} and \vec{H} . Assuming that the conductivity is zero (3.11) becomes

$$\frac{\bar{\epsilon}}{\Delta t} \left(\vec{E}^{+1}[0] - \vec{E}^0[0] \right) = \sum_{w=xyz} \bar{c}[+, w] \left(\vec{H}^{+\frac{1}{2}}[\delta_w] - \vec{H}^{+\frac{1}{2}}[-\delta_w] \right) \quad . \quad (3.22)$$

Note that $\bar{c}[+, w] = -\bar{c}[-, w]$ is an identity that comes from the antisymmetry of the curl. Inserting (3.20) and (3.21) into (3.22) gives a lengthy sum of complex exponentials. However, there is a common factor that can be removed from both sides and Euler's formula reduces the central differences into sine terms:

$$\frac{\bar{\epsilon}}{\Delta t} \vec{E}_0 \sin \left(\frac{\omega \Delta t}{2} \right) = \sum_{w=xyz} \bar{c}[+, w] \vec{H}_0 \sin(k_w \Delta w) \quad . \quad (3.23)$$

Similarly, the plane wave relation from the magnetic field update is:

$$-\frac{\bar{\mu}}{\Delta t} \vec{H}_0 \sin \left(\frac{\omega \Delta t}{2} \right) = \sum_{w=xyz} \bar{\gamma}[+, w] \vec{E}_0 \sin(k_w \Delta w) \quad . \quad (3.24)$$

Combining them gives an eigenvalue equation that yields the dispersion relation:

$$\sin^2 \left(\frac{\omega \Delta t}{2} \right) \vec{E}_0 = - \left(\sum_{w=xyz} \sum_{s=xyz} \Delta t^2 \bar{\epsilon}^{-1} \bar{c}[+, w] \bar{\mu}^{-1} \bar{\gamma}[+, s] \sin(k_w \Delta w) \sin(k_s \Delta s) \right) \vec{E}_0. \quad (3.25)$$

The simulation will remain stable if (3.25) predicts real frequencies. This is guaranteed as long as the eigenvalues of the matrix on the right-hand-side are less than or equal to one. Notice that (3.25) depends on direction via the sine terms on the right-hand-side. It can be shown that for a biaxial dielectric or magnetic material the

maximum value of the eigenvalues occurs when those sine terms are either ± 1 . This means that the axis of instability is still the grid diagonal, but each of the diagonals must be checked for stability since they may not share the same phase velocity. In summary, the maximum timestep must be computed from the largest eigenvalue of the total matrix along each of the 8 grid diagonals:

$$\Delta t_{\max} = \left(\max \left| \text{eig} \left(- \sum_{w=xyz} \sum_{s=xyz} \bar{\epsilon}^{-1} \bar{c}[+, w] \bar{\mu}^{-1} \bar{\gamma}[+, s] \text{dir}(w) \text{dir}(s) \right) \right| \right)^{-\frac{1}{2}}. \quad (3.26)$$

Note that $\text{dir}(w) = \pm 1$. For example, when considering the diagonal with unit vector $\hat{r} = (1, -1, 1)/\sqrt{3}$: $\text{dir}(x) = 1$, $\text{dir}(y) = -1$, and $\text{dir}(z) = 1$. Simulations are usually run with a slightly smaller timestep than Δt_{\max} . The constant of proportionality is commonly referred to as the Courant number ($S = \Delta t / \Delta t_{\max}$).

This formula was tested numerically for a variety of dielectric and magnetic materials in arbitrary orientations. A small domain ($7 \times 9 \times 11$) with unequal aspect ratio, periodic boundary conditions and several sources was setup and run for a long period of time. The electric field was excited with unit amplitude using a pulse that had significant bandwidth out to $1/10^{\text{th}}$ of the Nyquist limit. If the magnitude of any field in the domain exceeded 10^{30} within the first 65536 timesteps then that simulation run was considered unstable. Equation (3.26) was used to calculate the maximum timestep for each material. Then, the simulation was run starting with a Courant number of 1.05. It was decreased by 0.01 until the simulation was stable. The results of these runs for different materials in different orientations is given in Table 3.1 and show that the stability limit holds as predicted.

The dispersion relation is found by assuming that the waveform is well sampled

$\text{diag}(\bar{\epsilon})$	$\text{diag}(\bar{\mu})$	θ_z	θ_y	$\Delta t_{\text{theoretical}}$	$\Delta t_{\text{measured}}$	S
(1, 1, 1)	(1, 1, 1)	0°	0°	2.992 ps	2.992 ps	1.00
		36°	0°	2.992 ps	2.992 ps	1.00
		36°	11°	2.992 ps	2.992 ps	1.00
(4, 3, 3)	(1, 1, 1)	0°	0°	5.183 ps	5.183 ps	1.00
		36°	0°	5.183 ps	5.183 ps	1.00
		36°	11°	5.183 ps	5.183 ps	1.00
(1, 1, 1)	(3, 2, 2)	0°	0°	4.232 ps	4.232 ps	1.00
		36°	0°	4.232 ps	4.232 ps	1.00
		36°	11°	4.232 ps	4.232 ps	1.00
(4, 4, 2)	(1, 1, 1)	0°	0°	4.670 ps	4.670 ps	1.00
		36°	0°	4.670 ps	4.670 ps	1.00
		36°	11°	4.470 ps	4.470 ps	1.00
(1, 1, 1)	(3, 3, 2)	0°	0°	4.509 ps	4.509 ps	1.00
		36°	0°	4.509 ps	4.509 ps	1.00
		36°	11°	4.386 ps	4.386 ps	1.00
(4, 3, 3)	(3, 2, 2)	0°	0°	7.862 ps	7.862 ps	1.00
		36°	0°	7.687 ps	7.687 ps	1.00
		36°	11°	7.545 ps	7.545 ps	1.00
(3, 4, 5)	(1, 1, 1)	0°	0°	5.627 ps	5.684 ps	1.01
		36°	0°	5.261 ps	5.261 ps	1.00
		36°	11°	5.221 ps	5.221 ps	1.00
(1, 1, 1)	(2, 3, 4)	0°	0°	4.724 ps	4.771 ps	1.01
		36°	0°	4.316 ps	4.359 ps	1.01
		36°	11°	4.272 ps	4.272 ps	1.00
(3, 4, 5)	(2, 3, 4)	0°	0°	9.869 ps	9.869 ps	1.00
		36°	0°	9.216 ps	9.216 ps	1.00
		36°	11°	8.970 ps	8.970 ps	1.00

Table 3.1: Stability limit testing for lossless anisotropic materials in an infinite domain. Each simulation was run in a $7 \times 9 \times 11$ domain with a cell size of $1.3 \times 2.2 \times 1.5$ mm and periodic boundary conditions. A simulation is considered stable if it runs for 65536 timesteps and none of the fields within the domain have a magnitude exceeding 10^{30} . Each material was simulated with its principal axes aligned with the grid axes, rotated by 36° about the z -axis, and further rotated by 11° about the y -axis. The analytical stability limit was calculated using (3.26). The simulated limit was found by running a simulation at $S = 1.05$ and then decreasing it by 0.01 until the simulation was stable. Note that several simulations were measured as stable for $S = 1.01$. Further investigation showed that these simulations were unstable but the exponential increase was too slow to reach the assigned threshold within the simulation time.

in time:

$$\sin^2 \left(\frac{\omega \Delta t}{2} \right) \approx \left(\frac{\omega \Delta t}{2} \right)^2 . \quad (3.27)$$

This assumption is justifiable because the stability limit forces FDTD to sample better in time than in space. Assuming equal aspect ratio, $\Delta x = \Delta y = \Delta z = h$, the number of points per wavelength (N_λ) is defined to be

$$N_\lambda = \frac{2\pi}{kh} . \quad (3.28)$$

Equation (3.25) is then rearranged and the numerical phase velocity (c^*) is found from the eigenvalues of a matrix:

$$(c^*)^2 = - \left(\frac{c_0 N_\lambda}{\pi} \right)^2 \text{eig} \left(\sum_{w=xyz} \sum_{s=xyz} \bar{\epsilon}_r^{-1} \bar{c}[+, w] \bar{\mu}_r^{-1} \bar{\gamma}[+, s] \sin \left(\frac{r_w \pi}{N_\lambda} \right) \sin \left(\frac{r_s \pi}{N_\lambda} \right) \right) . \quad (3.29)$$

The discussion concerning the existence of the eigenvalues of this equation is the same as the one given in Section 2.2.

Although the form of (3.29) is not very enlightening it is very easy to solve numerically once the material parameters and direction of propagation have been decided. It can also be shown that this form reduces to the same dispersion relation as a Yee grid with equal cell size if it is assumed that the materials are isotropic. That result was also proven by Garcia [2]. Figures 3.3 to 3.5 plot the normalized numerical phase velocity error ($|c^* - c_0|/c_0$) as a function of direction. They show that even though the error depends on the polarization, the dependence on direction takes the same form as in the isotropic case. That is, the error is maximum in the cardinal directions and minimum along the grid diagonals.




Figure 3.3: Normalized numerical phase velocity error as a function of direction on a Lebedev grid for an isotropic dielectric. The amount of error depends on the number of grid points per wavelength causing dispersion that is worst along the cardinal axes.




Figure 3.4: Normalized numerical phase velocity error as a function of direction on a Lebedev grid for uniaxial dielectric. The error depends only slightly on the polarization and the magnitude of the error is comparable to isotropic materials.




Figure 3.5: Normalized numerical phase velocity error as a function of direction on a Lebedev grid for biaxial dielectric. The grids cardinal axes exhibit the worst dispersion and the stability limiting axis is the one of the grid diagonals despite phase velocity depending on polarization and direction of propagation.

3.5 Equal Memory Comparison

In order to understand the practicality of the Lebedev grid, it must be compared to existing Yee grid methods. However, making a fair comparison is a difficult task because the two grids excel in different areas. This section discusses several different aspects that significantly affect the practicality of a computational scheme. Things to consider encompass details from the entire thesis. Therefore, it is recommended that the reader return to this section after reading Chapter 4.

There are four important aspects that determine the feasibility of a computational method. First, the method must be accurate enough to replace laboratory measurements during the design process. In FDTD, the two potential limiting factors are available time and memory. The computational grid must be small enough to fit in the random access memory (RAM) of the computer. Then, the time required for the

simulation to complete is largely determined by how many calls to RAM are required during each update. Finally, the method should be simple enough that complex circuit components can be simulated with minimal programming effort. In summary, the four factors to consider are numerical accuracy, total required memory, total calls to memory, and program simplicity.

Consider the Yee grid in lossless anisotropic electric and magnetic media. In this case, \vec{D} and \vec{B} must be collocated with \vec{E} and \vec{H} and the total memory requirement of the Lebedev grid is twice that of the Yee grid. In isotropic media \vec{D} and \vec{B} need not be stored and the Lebedev grid uses four times as much memory. However, in lossy anisotropic media, the losses must be extrapolated using field values from the previous two timesteps and the Lebedev grid uses one third more memory. For simplicity, the intermediate case will be used and the Yee grid will be considered to use half the memory of the Lebedev grid. Therefore, a cell size on the Yee grid of $\Delta z_{\text{Yee}} = \Delta z_{\text{Leb}} / \sqrt[3]{2} \approx \Delta z_{\text{Leb}} / 1.26$ discretizes the same volume with the same amount of memory.

In lossless anisotropic electric and magnetic media the Lebedev grid accesses the three components at the gridpoint and two components at its six nearest neighbors. There are 4 \vec{E} gridpoints and 4 \vec{H} gridpoints per unit cell. Therefore, each timestep and grid cell requires 120 calls to memory. On the Yee grid the update is performed in four stages. Each component is updated separately and 5 memory accesses are required to evaluate the curl equations and 9 to enforce the constitutive relations. There are 3 components of each field type per grid cell resulting in 84 memory accesses per cell per timestep. However, on an equal memory basis the number of grid cells doubles and the timestep decreases. The stability limit on both grid takes the same

form ($\Delta t = \Delta z/c$). Therefore, the ratio of required memory accesses on the Lebedev grid to the Yee grid is $120:168\sqrt[3]{2}$ or 1:1.4. Furthermore, the grid storage format of the Lebedev grid allows better caching of the memory accesses (see Appendix C for more details).

In summary, the Lebedev grid has a shorter compute time when the computational domain is discretized with the same amount of memory. Therefore, comparing accuracy on an equal memory basis is a more stringent requirement than comparing on the basis of equal compute time. The rest of the thesis will show that the programming structure is simpler on the Lebedev grid in all situations but the numerical results are more accurate only in certain situations. Therefore, the applicability of the Lebedev grid is situation dependent and understanding the trade-offs will help engineers decide on the appropriate grid for a given problem. This discussion is continued throughout the course of the thesis.

3.6 Grid Degeneracy

It has been shown that the algorithm is stable, but there is still a major issue that needs to be addressed. If a single grid point is excited in an isotropic medium and the simulation is marched forward in time, only some of the components at each gridpoint will attain non-zero values (Figure 3.6). It turns out, that the Lebedev grid can be decomposed into four Yee grids (Figure 3.7) [5] and it can be shown that these grids do not couple to one another in isotropic media [25]. The solution to this problem is to excite and sample from all four Yee grids. That way, each Yee grid runs its own isotropic solution, and the perturbation of anisotropy is added from the isotropic solutions on the surrounding grids. Furthermore, boundary conditions align




Figure 3.6: Magnetic field vector plots of the initial propagation of a pulse on a Lebedev grid with averaged sources and sampling (left) and without averaging (right). The arrows represent both x and y field components at each location because the Lebedev grid uses collocated components. The arrows in the unaveraged simulation point solely in the x and y -directions because only one of the Yee subgrids has been excited. This simulation was run in a vacuum about an electric dipole.

differently with each of the grids and boundary condition errors tend to cancel when the solutions are averaged together [5]. In order to derive the form of the excitation, a more intuitive derivation of the update equations is presented.

3.7 Integral Derivation of Update Equations

The integral derivation of the FDTD update equations on a Yee grid is well known [13]. The decomposition of the Lebedev grid into its four Yee grids gives the starting point for this derivation. The integral form of Maxwell's equations is solved by equating the flux through a surface integral to the line integral along the contour that defines




Figure 3.7: Decomposition of the Lebedev grid (center) into four offset Yee grids (surrounding). The Yee grids are displaced from one another by half the cell size in two of the x , y or z -directions.

the edge of the surface:

$$\oint \vec{H} \cdot d\vec{L} = \iint \left(\bar{\epsilon} \frac{\partial \vec{E}}{\partial t} + \bar{\sigma} \vec{E} + \vec{J} \right) d\vec{S} \quad , \quad (3.30)$$

$$\oint \vec{E} \cdot d\vec{L} = \iint \left(-\bar{\mu} \frac{\partial \vec{H}}{\partial t} - \bar{\sigma}^* \vec{H} - \vec{M} \right) d\vec{S} \quad . \quad (3.31)$$

The contour is chosen to be a square that lies in the xy , yz , or zx -plane centered about the gridpoint being updated (Figure 3.8). The normal component in the center and the components parallel to the lines each belong to the same underlying Yee grid. In isotropic media only these components are related to one another, but the tensors in the surface integral couple in all the components at the point being updated. It can be seen that the updates only depend on the tangential components of the neighboring




Figure 3.8: The finite difference equations can be derived by approximating the line integrals along the edges of square integration loops that surround a gridpoint by the midpoint rule. On the Lebedev grid (right) there are two overlapping sets of integration loops on any given plane whereas the integration loops cover the Yee grid (left) without any overlap or gaps. Any current through a plane on the Lebedev grid passes through more than one integration loop and therefore the excitation must be distributed on the Lebedev grid.

points. This fact can be used to optimize the code, but it is not generally true. One known exception is a material interface.

From here the derivation is straightforward. The line integrals are broken into the four sections:

$$\oint \vec{H} \cdot d\vec{L} = \int_{-\frac{\Delta x}{2}}^{\frac{\Delta x}{2}} H_x(x, -\frac{\Delta y}{2}, 0) dx + \int_{-\frac{\Delta y}{2}}^{\frac{\Delta y}{2}} H_y(\frac{\Delta x}{2}, y, 0) dy \\ + \int_{\frac{\Delta x}{2}}^{-\frac{\Delta x}{2}} H_x(x, \frac{\Delta y}{2}, 0) dx + \int_{\frac{\Delta y}{2}}^{-\frac{\Delta y}{2}} H_y(-\frac{\Delta x}{2}, y, 0) dy \quad , \quad (3.32)$$

and each section is evaluated by the midpoint rule:

$$\oint \vec{H} \cdot d\vec{L} \approx \Delta x H_x[0, -1, 0] + \Delta y H_y[1, 0, 0] \\ - \Delta x H_x[0, 1, 0] - \Delta y H_y[-1, 0, 0] \quad . \quad (3.33)$$

This approximation is the same as the one obtained using the curl equations other than a factor of $\Delta x \Delta y$:

$$\frac{\oint \vec{H} \cdot d\vec{L}}{\Delta x \Delta y} = \left(\nabla \times \vec{H} \right)_z . \quad (3.34)$$

Each of the electric field surface integrals are approximated by a two dimensional midpoint rule using the grid point at the center of the loops:

$$\begin{aligned} \iint \bar{\epsilon} \frac{\partial \vec{E}}{\partial t} \cdot d\vec{S} &= \int_{-\frac{\Delta x}{2}}^{\frac{\Delta x}{2}} \int_{-\frac{\Delta y}{2}}^{\frac{\Delta y}{2}} \left(\epsilon_{zx} \frac{\partial E_x}{\partial t} + \epsilon_{zy} \frac{\partial E_y}{\partial t} + \epsilon_{zz} \frac{\partial E_z}{\partial t} \right) dy dx \\ &\approx \Delta x \Delta y \left(\epsilon_{zx} \frac{\partial E_x}{\partial t}[0] + \epsilon_{zy} \frac{\partial E_y}{\partial t}[0] + \epsilon_{zz} \frac{\partial E_z}{\partial t}[0] \right) . \end{aligned} \quad (3.35)$$

If the time derivatives are evaluated with central differences then the result can be rearranged to the previous scheme given by (3.12) and (3.16). The key difference is that the current terms can be taken into account. The surface integrals containing current terms should be evaluated directly. They should not be evaluated with the midpoint rule. Let I_z be the imposed current passing through an integration loop in the xy -plane:

$$I_z = \int_{-\frac{\Delta x}{2}}^{\frac{\Delta x}{2}} \int_{-\frac{\Delta y}{2}}^{\frac{\Delta y}{2}} J_z(x, y, 0) dy dx . \quad (3.36)$$

The imposed current (\vec{I}) may be in any direction. The other two formulas are found by circulating the components $x \rightarrow y \rightarrow z \rightarrow x$. The algebraic manipulations introduce a factor that normalizes the area of the integration loops. That is, let $\bar{\bar{A}}$ be the diagonal matrix defined by

$$A_{uv} = \frac{\delta_{uv} \Delta u}{\Delta x \Delta y \Delta z} . \quad (3.37)$$

Then, the source driven update equation is

$$\bar{\bar{a}}\vec{E}^{+1}[0] = \bar{\bar{b}}\vec{E}^0[0] + \sum_{p=\pm 1} \sum_{w=xyz} \bar{\bar{c}}[p, w] \vec{H}^{+\frac{1}{2}}[p\delta_w] + \bar{\bar{A}}\vec{I}^{+\frac{1}{2}} \quad . \quad (3.38)$$

The sourced update for the magnetic field is found using duality:

$$\bar{\bar{\alpha}}\vec{H}^{+\frac{1}{2}}[0] = \bar{\bar{\beta}}\vec{H}^{-\frac{1}{2}}[0] + \sum_{p=\pm 1} \sum_{w=xyz} \bar{\bar{\gamma}}[p, w] \vec{E}^0[p\delta_w] + \bar{\bar{A}}\vec{I}^{*0} \quad . \quad (3.39)$$

It is important to notice that the integration loops in any given plane overlap with one another (Figure 3.8). This means that current will always pass through more than one loop and that boundaries may align differently with different loops. Rigorously considering these different cases is the key to maintaining the same isotropic solution on the four Yee grids. When integration loops are treated differently, different underlying Yee grids approximate different circuits. This leads to spurious solutions, especially since anisotropy couples the four solutions together.

3.8 Sources and Sampling

The previous section showed that a current source passes through more than one integration loop. This must be taken into account in order to set up an appropriate excitation. Assume that a line current source has been placed along the z -axis at the location $[i, j, k]$, where i and j are fixed. There are three types of electric field points surrounding the source. The first type of electric field points share locations with the

source itself. Remembering that the flux passes through two integration loops gives

$$J_z[i, j, k] = \frac{J_0}{2} \quad . \quad (3.40)$$

The second type of electric field points are at the locations $[i \pm 1, j \pm 1, k]$. For these loops, the source is on the corner of the integration loop. Therefore the flux through these loops is a quarter of the flux through a loop centered on the source:

$$J_z[i \pm 1, j \pm 1, k] = \frac{J_0}{8} \quad . \quad (3.41)$$

The planes above and below the first two types of electric field points are located at $[i \pm 1, j, k]$ or $[i, j \pm 1, k]$. The line source intersects these planes on the very edge of the integration loop. It follows that the flux is divided by two for the ratio of areas and two again because of the overlapping integration loops:

$$J_z[i \pm 1, j, k] = \frac{J_0}{4} \quad , \quad J_z[i, j \pm 1, k] = \frac{J_0}{4} \quad . \quad (3.42)$$

Figure 3.9 shows the relative excitation weight of the points adjacent to the source.

An infinitesimal dipole can be considered a short line source located at $[i, j, k]$. In order to reproduce the stencil derived in [5] that was based on exciting all grids in the same way, the line source has half its weight in the plane containing the point and a quarter of its weight in the planes above and below the dipole. The line source aligns with one point directly. The weight of the source in that plane is $J_0/2$ and there are




Figure 3.9: Source weighting of a line of current traveling in the z -direction and passing through the center points.

two sets of integration loops in that plane:

$$J_z[i, j, k] = \frac{J_0}{4} \quad . \quad (3.43)$$

In the same plane, the source aligns with the corners of the integration loop for the points at $[i \pm 1, j, k]$:

$$J_z[i \pm 1, j \pm 1, k] = \frac{J_0}{16} \quad . \quad (3.44)$$

Above and below the plane of the source the current aligns with the edges of the integration loops:

$$J_z[i \pm 1, j, k \pm 1] = \frac{J_0}{16} \quad , \quad J_z[i, j \pm 1, k \pm 1] = \frac{J_0}{16} \quad . \quad (3.45)$$

The total result is that the point being excited has a weight of $I/4$ and the surrounding




Figure 3.10: Source weighting of an electric dipole located in the center of the grid cell. The weighting does not depend on the orientation of the dipole.

points have a weight of $I/16$. The dipole excitation can be oriented in any direction because of the symmetry of the weights. These weights are shown in Figure 3.10.

A plane source is assumed to be a sheet of current that lies in a plane of points. In this case, the current passes directly through the integration loops that contain the points and passes along the edge of the integration loops of the points one half cell away from the source plane. Therefore, the points in the plane have a weight of $J_0/2$ and the points in the neighboring planes have a weight of $J_0/4$. These weights are shown in Figure 3.11.

The discussion in this section was focused on setting up the grid excitation. However, there is also ambiguity with how to sample data from the grid. For example, the power traveling down a transmission line is measured by integrating the flux of the Poynting vector through a surface. The fields at a gridpoint represent the amount of flux passing through the integration loop from Section 3.7. Since these loops overlap,




Figure 3.11: Source weighting about a sheet of current lying in the yz -plane and passing through the center points.

the actual flux must be some average of the contributing points. On the other hand, if the electric field is being measured then the value at a single gridpoint should be used. Ideally, it does not matter which Yee grid each component belongs to because all four grids propagate the exact same solution. That said, experience shows that averaging the results from all four Yee grids reduces errors caused by the approximate excitation. Davydycheva argues that the different representation of boundary conditions on each grid leads to reduced error only when the average solution is used [5]. Therefore, every sampled electric field was measured with a quarter of its weight at the gridpoint being considered and a sixteenth of its weight at the nearest 12 electric field gridpoints (Same weights as infinitesimal dipole excitation).




Figure 3.12: Setup of computational domain for simulating plane waves using only equations derived for an infinite domain. All boundaries are periodic and the excitation waveform is the same at each point so that a plane wave is launched in the $\pm z$ -directions. A short line is sampled over time and the simulation is run only long enough for a single pulse to travel through the sampled line (See Figure 3.13).

3.9 Plane Waves in an Anisotropic Material

It is now possible to put the discussed tools together and setup a simulation to compare numerical plane waves to the analytical solutions. To do this, an infinite array of dipoles is simulated. The dipole lies in the center of the xy -plane and launches plane waves in the z -direction for all frequencies. At higher frequencies, plane waves are also launched in other directions depending on the ratio between wavelength and the spacing between the dipoles. The advantage of this setup is that it uses only periodic boundary conditions and the update equations in a bulk material. Therefore, the stability criterion and dispersion relation can be investigated without worry of components in the domain affecting the results.

The dipole array is setup using the elongated domain shown in Figure 3.12. The computational domain is extremely long in the z -direction and relatively short in the x and y -directions. There is a dipole source lying in the center of the excitation plane. This single dipole source is equivalent to an array of dipoles because of the

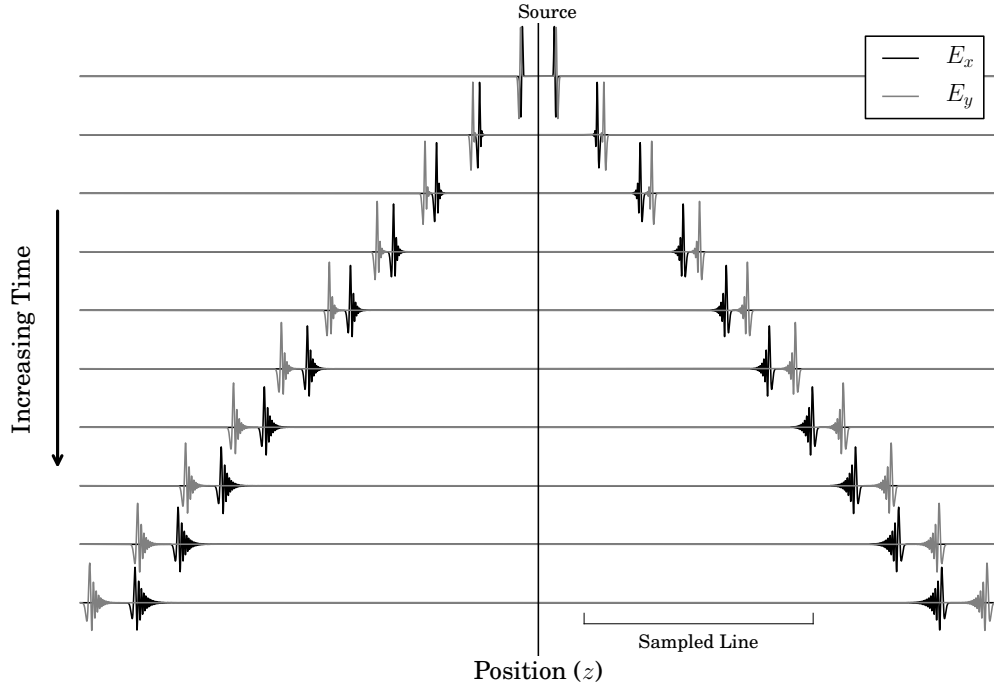


Figure 3.13: Plane wave pulses launched from a uniform plane source in aligned sapphire. The pulses of the two polarizations split because they have different propagation constants. A line of samples is measured over time and transformed to frequency and wavenumber to extract the propagation constant as a function of frequency for each mode (See Figure 3.14).

periodicity in the x and y -directions. The z -direction also has periodic boundary conditions but they are not used. A pulse is launched from the center of the domain and is propagated until it reaches the end of the computational domain. A line of samples records the electric fields over t and z (Figure 3.13). In a post-processing stage this data is transformed to frequency (f) and the propagating constant (k_z) (Figure 3.14) and used to extract the k_z vs f characteristic for each propagating mode as described in Appendix B.1.

The dimensions of the domain were chosen to conform to the other simulations in the thesis. The number of cells in the domain was $14.5 \times 8.5 \times 4750.5$ (half integer

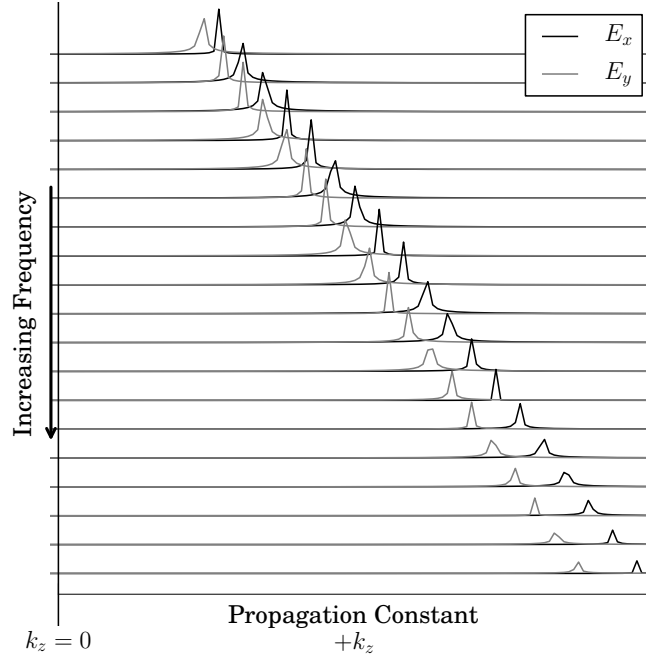


Figure 3.14: Fourier transform in time and space of sampled lines from Figure 3.13. The propagation constant as a function of frequency is found for each polarization by zooming in on the ridgeline with the chirp Z-transform at fixed frequencies. Actual results presented contain more curves because the uniform plane wave source was changed to a dipole in the center of the plane (Figure 3.15).

cells for periodic boundary conditions). An equal aspect ratio was used with a cell spacing of $\Delta x = \Delta y = \Delta z = h = 0.814$ mm. The plane wave source was the first derivative of a Gaussian and had bandwidth up to 14 GHz as defined in Appendix A. Each simulation was run with a Courant number of 0.99. An equivalent Yee grid simulation was set up using an equal memory comparison. The Yee grid spacing was $\Delta z_{\text{Yee}} = \Delta z_{\text{Leb}} / \sqrt[3]{2}$ because the Lebedev grid uses twice as much memory to discretize the same 3D volume.

The numerical results are compared with an analytical solution. The analytical solution for the dipole array is the same as that of a diffraction grating. For a

diffraction grating, the location of the observed maxima corresponds to the direction of the plane wave launched from the dipole array. The allowed directions are a function of wavelength (λ) and the spacing between dipoles (L_x, L_y).

The exact formula for the direction of propagation: $\hat{r} = (\pm \sin(\theta_x), 0, \cos(\theta_x))$, is given by

$$\sin(\theta_x) = \frac{m_x \lambda}{L_x} \quad , \quad m_x \in \mathbb{Z} \quad . \quad (3.46)$$

An equivalent formula holds for the y -direction. In anisotropic material this formula is more complicated because the wavelength depends on polarization and direction of propagation. The analytical solution is found by iteratively solving for the plane wave wavelength using the method in Section 2.2 ($\lambda = c_0/nf$) and resolving for the angle of maximum interference (θ_x).

Figure 3.15 shows the propagation constants in the z -direction ($\beta = k \cos(\theta_x)$) for the three materials given in Section 2.3. The plot shows that the Lebedev and Yee grid based results both agree well with the analytical solution. The largest error is at high frequency. For poorly sampled wavelengths the numerical phase velocity error causes the numerical k_z to lie above the analytical k_z . The Lebedev grid shows greater error than the Yee grid because the Yee grid uses a smaller cell size to give an equal memory comparison. These results suggest that when simulating homogeneous materials the interpolation error on the Yee grid is not significant enough to merit the extra memory required by the Lebedev grid. Note that this evaluation is based solely on accuracy and memory. Total compute time and program simplicity are ignored.

It is interesting to note the effect of anisotropy on the interference pattern. In the isotropic case there are only three observed waves. The $m = 0$ wave at all frequencies corresponding to a plane wave in the z -direction. A plane wave for $m_x = 1$ that is




Figure 3.15: Propagation constant versus frequency for plane waves launched from an array of dipoles in three different media simulated using the Lebedev grid (solid) and the Yee grid (dashed). The data is obtained from a single simulation by Fourier transforming a line of electric field samples as described in Appendix B.1. The numerical solutions differ from the analytical solution (dotted) the least when sampled at more than 10 points per wavelength (N_λ).

launched when $\lambda < L_x$ and a similar plane wave for $m_y = 1$. When anisotropy is introduced there are two plane waves for $m_x = 1$, one for each polarization. When the anisotropy is misaligned then a third plane wave is visible. This occurs because the wavelengths in the $\pm x$ -direction are different for the extraordinary wave but the same for the ordinary wave.

Figure 3.16 shows the basis linear polarizations in anisotropic media for the plane wave launched along the z -axis. The image was generated by plotting the transverse components of the electric field over x and y at a specified k_z and f . In the isotropic case the polarizations are degenerate so the plotted polarizations for both modes are the same as the source excitation. However, in anisotropic media the propagation constants differ and the two polarization states can be distinguished. This type of plot is not particularly useful for plane waves, but provides valuable information about transmission lines when used in more complicated simulations.

3.10 Conservation of Charge

The Yee grid is renowned for providing a highly convergent and physically relevant approximation to Maxwell's equations. For example, the grid satisfies the divergence equations because the grid conserves charge [13]. This can be shown by proving that the numerical scheme has a consistent calculus [26]. That is, by showing that the divergence of the curl is zero.

Consider the continuity equation. It is derived by taking the divergence of the curl equation, (2.2), and substituting in the divergence equation, (2.6):

$$\frac{\partial \rho_e}{\partial t} = \nabla \cdot (\nabla \times \vec{H}) - \nabla \cdot \vec{J}_{\text{total}} \quad . \quad (3.47)$$

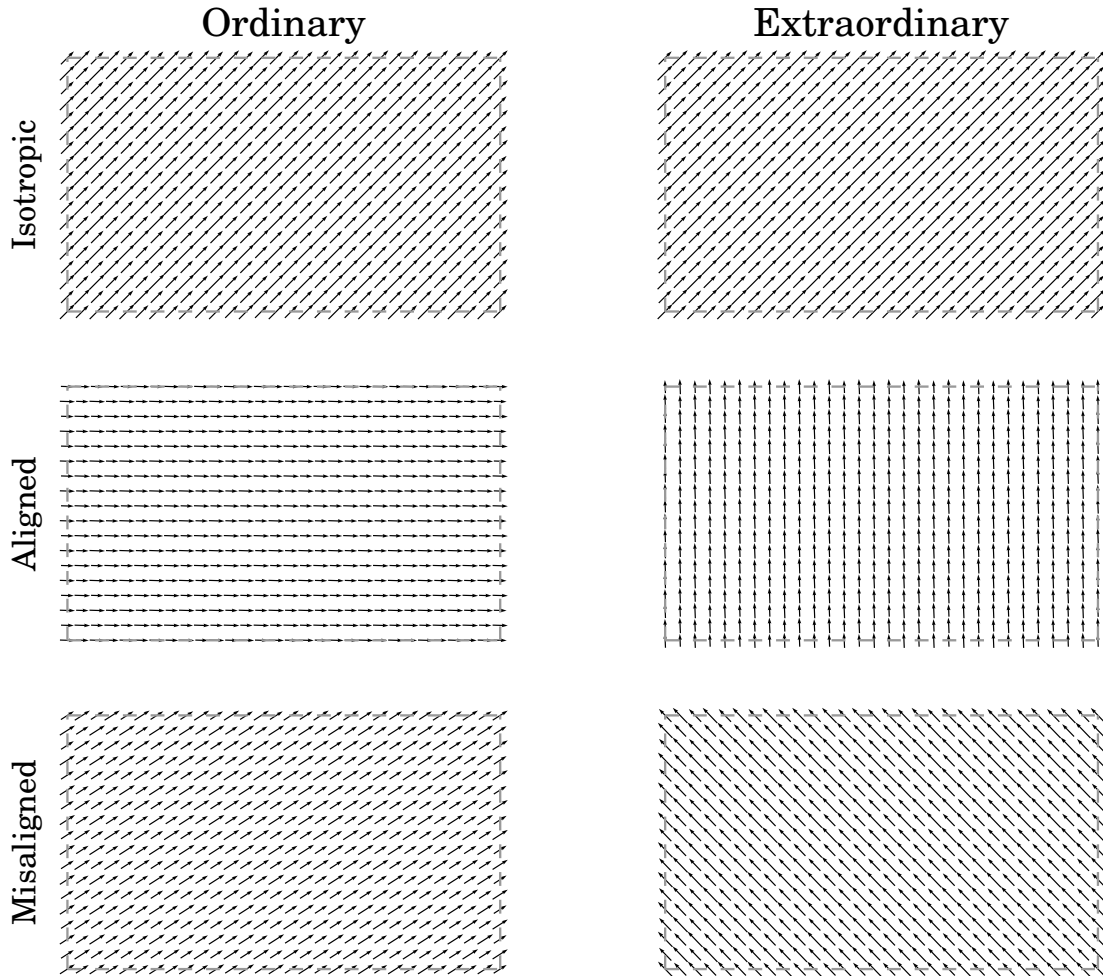


Figure 3.16: Transverse components of the electric field for a plane wave traveling in three different media. The polarization shown in both isotropic plots is the polarization of the excitation because all polarizations have the same spatial dependence and cannot be distinguished. In anisotropic dielectrics only two linear polarizations are allowed. Their orientation depends on the direction of propagation with respect to the material's optical axis. The plots were obtained using the method in Appendix B.2.

It is apparent that charge is conserved in a source free region when the divergence of the curl is zero. This is an identity of continuous calculus but is not necessarily satisfied in discrete calculus.

The discrete approximations to the curl and the divergence are

$$\nabla \times \vec{E}[0] = \sum_{p=\pm 1} \sum_{w=xyz} \bar{c}[p, w] \vec{E}[p\delta_w] \quad , \quad c_{uv}[p, w] = \varepsilon_{uvw} \frac{p}{\Delta w} \quad , \quad (3.48)$$

and

$$\nabla \cdot \vec{E}[0] = \sum_{p=\pm 1} \sum_{w=xyz} \bar{d}[p, w] \vec{E}[p\delta_w] \quad , \quad d_{uv}[p, w] = \delta_{uw} \delta_{vw} \frac{p}{\Delta w} \quad . \quad (3.49)$$

It follows that the divergence of the curl is

$$\nabla \cdot (\nabla \times \vec{E}) = \sum_{p=\pm 1} \sum_{w=xyz} \sum_{n=\pm 1} \sum_{s=xyz} \bar{d}[p, w] \bar{c}[n, s] \vec{E}[p\delta_w + n\delta_s] \quad . \quad (3.50)$$

Although this equation looks hopeless, each matrix is actually extremely sparse. In fact, $\bar{d}[p, w]$ has one non-zero entry and $\bar{c}[n, s]$ has two non-zero entries. More importantly, (3.50) shows that the divergence of the curl depends on the point at the location being considered and the points that are two steps away. Consider the case where $w = s$. For example, if $w = s = x$, then

$$\bar{d}[p, x] \bar{c}[n, x] = \begin{pmatrix} \frac{p}{\Delta x} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & -\frac{n}{\Delta x} \\ 0 & \frac{n}{\Delta x} & 0 \end{pmatrix} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad . \quad (3.51)$$

independently of p and n . A similar equation holds for $w = s = y$ and $w = s = z$.

Consider now the case where $w \neq s$ by letting $w = x$ and $s = y$:

$$\bar{d}[p, x] \bar{c}[n, y] = \begin{pmatrix} \frac{p}{\Delta x} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} 0 & 0 & \frac{n}{\Delta y} \\ 0 & 0 & 0 \\ -\frac{n}{\Delta y} & 0 & 0 \end{pmatrix} = \begin{pmatrix} 0 & 0 & \frac{pn}{\Delta x \Delta y} \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} . \quad (3.52)$$

It follows that the $E_z[n, p, 0]$ is brought into the sum in (3.50). However, all the grid points with $w \neq s$ are brought into the sum twice. For the case being considered, the second term is when $w = y$ and $s = x$:

$$\bar{d}[p, y] \bar{c}[n, x] = \begin{pmatrix} 0 & 0 & 0 \\ 0 & \frac{p}{\Delta y} & 0 \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & -\frac{n}{\Delta x} \\ 0 & \frac{n}{\Delta x} & 0 \end{pmatrix} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & -\frac{pn}{\Delta y \Delta x} \\ 0 & 0 & 0 \end{pmatrix} . \quad (3.53)$$

$E_z[n, p, 0]$ is once again brought into the sum, but this time with opposite sign. Therefore, all of the terms with $w = s$ are immediately zero and all of the terms with $w \neq s$ come in pairs that introduce the same component with equal magnitude but opposite sign. Summing all the pairs together cancels all the non-zero terms proving that the Lebedev grid has a consistent calculus:

$$\nabla \cdot (\nabla \times \vec{E}) = 0 . \quad (3.54)$$

3.11 Infinite Domain Summary

In this chapter the Lebedev grid was introduced as the natural extension to the Yee grid for anisotropy because it is based on central differences. The update equation

was cast in a form that is easy to program. This form will be maintained as boundary conditions are introduced. A proof of stability was given for lossless dielectric and magnetic materials. A series of numerical tests showed that simulations exhibit exponential instability at 1% above the calculated stability limit and are stable within the limit. The dispersion relation was derived and plots for different materials show that the dispersion error follows the same trends as on the Yee grid. That is, the numerical phase velocity depends on direction, it is smaller than the analytical phase velocity, and the discrepancy is greatest along the cardinal axes. It was shown that the Lebedev grid is composed of four Yee grids and that to remove spurious solutions approximations must be equivalent on each Yee grid. An integral method of deriving the update equations that provides insight to the nature of the spurious solutions was provided. This derivation was used to motivate distributed excitations and averaged sampling. Plane waves were simulated in three anisotropic dielectric materials. The results show that the computational error in plane wave simulations is dominated by the numerical phase velocity error and therefore the Yee grid produces better results based on an equal memory comparison. Finally, it was proven that the Lebedev grid has a consistent calculus. It follows that the divergence equations are satisfied by the method even though the method is derived from the curl equations. This chapter laid a rigorous foundation for simulating anisotropic media. The next step is to introduce some basic boundary conditions and simulate actual microwave transmission lines.

Chapter 4

Planar Boundary Conditions and Material Interfaces

This chapter introduces several boundary conditions that are fundamental to FDTD transmission line and antenna analysis. Specifically, perfect conductors, material interfaces, and absorbing boundaries that align with a rectangular computational domain are considered assuming anisotropic constitutive relations. In all cases, the goals stressed throughout the thesis are maintained. Namely, the algorithm should be easy to implement, parallelizable, stable, and not excite spurious solutions.

4.1 Convolutional Perfectly Matched Layers

Absorbing boundary conditions are used when a finite computational domain is meant to extend to infinity. For example, when an entire pulse needs to pass through a line of samples it is inefficient to simulate a portion of the computational domain as was done in Section 3.9. Instead, a reflectionless boundary is placed at the ends of the sampled line so that any portion of the pulse that has passed the line need not be simulated any longer. Absorbing boundary conditions are similarly used in radiation and scattering problems.

The convolutional perfectly matched layer (CPML) [27] is the most successful

absorbing boundary condition for FDTD [13]. The CPML is based on the idea of reflectionless absorbing materials. The conductivity and permittivity are increased together such that the material still has the same wave impedance. The frequency dependence of the material properties is taken into account in the derivation and a convolution must be evaluated in the time domain. Fortunately, the convolution has a simple form that is well approximated by a recursive updating scheme [27]. Therefore, the CPML is implemented by storing and updating several extra variables in the CPML regions. The formulas and variables stored are the same for the Lebedev grid as they are for the Yee grid. Anisotropy is easily taken into account because the derivation for the CPML formulation is material independent [28].

4.1.1 Implementation

The update equation in a CPML region can be split into two parts. The auxiliary terms (ψ) that are recursively updated are added to the regular update equation. For example, the PML region around $x = 0$ is updated by

$$\bar{a}\vec{E}^{+1} = \bar{b}\vec{E}^0 + \sum_{p=\pm} \sum_{w=x,y,z} \bar{c}[p,w] \vec{H}^{+\frac{1}{2}} + \begin{pmatrix} 0 \\ -1 \\ 0 \end{pmatrix} \psi_{eyx}^{+\frac{1}{2}} + \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \psi_{ezx}^{+\frac{1}{2}} \quad , \quad (4.1)$$

$$\bar{\alpha}\vec{H}^{+\frac{1}{2}} = \bar{\beta}\vec{H}^{-\frac{1}{2}} + \sum_{p=\pm} \sum_{w=x,y,z} \bar{\gamma}[p,w] \vec{H}^0 + \begin{pmatrix} 0 \\ -1 \\ 0 \end{pmatrix} \psi_{hyx}^0 + \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \psi_{hzx}^0 \quad . \quad (4.2)$$

The regular update and the auxiliary terms are split from one another in the program. With this form, the entire FDTD domain is still updated in the same way. Then, the

auxiliary terms are added only in the CPML regions. Before the auxiliary terms are added they are updated from the nearby fields:

$$\psi_{eyx}^{+\frac{1}{2}} = b_{ex}\psi_{eyx}^{-\frac{1}{2}} + a_{ex}\frac{\partial H_z}{\partial x}^{+\frac{1}{2}}, \quad (4.3)$$

$$\psi_{ezx}^{+\frac{1}{2}} = b_{ex}\psi_{ezx}^{-\frac{1}{2}} + a_{ex}\frac{\partial H_y}{\partial x}^{+\frac{1}{2}}, \quad (4.4)$$

$$\psi_{hyx}^{+\frac{1}{2}} = b_{hx}\psi_{hyx}^{-\frac{1}{2}} + a_{hx}\frac{\partial E_z}{\partial x}^{+\frac{1}{2}}, \quad (4.5)$$

$$\psi_{hzx}^{+\frac{1}{2}} = b_{hx}\psi_{hzx}^{-\frac{1}{2}} + a_{hx}\frac{\partial E_y}{\partial x}^{+\frac{1}{2}}, \quad (4.6)$$

where

$$b_{ex} = e^{\left(\frac{\sigma_{pex}}{\kappa_{ex}} + \alpha_{pex}\right)}, \quad (4.7)$$

$$a_{ex} = \frac{\sigma_{pex}}{\sigma_{pex} + \alpha_{ex}\kappa_{ex}} (b_{ex} - 1). \quad (4.8)$$

The derivatives are evaluated using a central difference. However, a stretched coordinate system is typically used in the PML region. Therefore, the derivative is evaluated as

$$\frac{\partial H_y}{\partial x} = \frac{H_y[1, 0, 0] - H_y[-1, 0, 0]}{\kappa_{ex}\Delta x}.$$

Note that the stretched coordinate system must also be taken into account in the curl matrices. This is easily done by evaluating $\bar{c}[p, w]$ and $\bar{c}[p, w]$ with $\Delta x \rightarrow \kappa_{ex}\Delta x$.

The constants σ , κ , and α are chosen by the user. They are graded throughout the CPML region in response to a trade-off between reflection off the domain-CPML interface and reflection off the domain boundary after propagating through the CPML region [13]. In this thesis, a polynomial grading scheme was used and the values of

the parameters were calculated from the effective isotropic material parameters using:

$$m = 3 \quad , \quad (4.9)$$

$$m_\alpha = 2 \quad (4.10)$$

$$\sigma_{\max} = 0.6 \frac{m+1}{\Delta x} \sqrt{\frac{\epsilon}{\mu}} \quad , \quad (4.11)$$

$$\kappa_{\max} = 11 \quad , \quad (4.12)$$

$$\alpha_{\max} = 0.0 \quad , \quad (4.13)$$

$$N_{\text{PML}} = 10 \quad , \quad (4.14)$$

$$t[i] = \frac{i + 1/2}{2N_{\text{PML}} + 1/2} \quad , \quad (4.15)$$

$$\kappa[i] = 1 + t[i]^m (\kappa_{\max} - 1) \quad , \quad (4.16)$$

$$\sigma[i] = t[i]^m \sigma_{\max} \quad , \quad (4.17)$$

$$\alpha[i] = (1 - t[i])^{m_\alpha} \alpha_{\max} \quad , \quad (4.18)$$

where i is the index that is zero at the CPML-domain boundary and $2N_{\text{PML}}$ at the edge of the computational domain. Motivations for the chosen parameters are taken from a combination of [13, 27, 28, 29].

4.1.2 Measured Reflection Coefficient

The CPML formulation was verified by measuring the reflection coefficient of plane waves normally incident on a the CPML using the setup in Figure 4.1. The reflection coefficient of the plane wave simulations using periodic boundaries and a windowed line of samples (Section 3.9) was also measured and shows the noise floor of the reflection coefficient calculation. Figure 4.2 shows that at low frequencies the CPML




Figure 4.1: Setup of computational domain for measuring the reflection coefficient of a CPML. A pulse is propagated until the reflection off the CPML at $z = L_z$ has passed the excitation plane. The CPML at $z = 0$ is three times larger than the other and is considered ideal. The reflection coefficient of each mode may be found as a function of frequency as described in Appendix B.3.

reflection coefficient is at the noise floor of the method. At higher frequencies the CPML error is above the noise floor of the reflection coefficient method, but the reflection coefficient remains below -40 dB. For all frequencies, polarizations, and materials, the measured reflection coefficient is negligible compared to any measurements made in other parts of the thesis. Therefore, the CPML absorbing boundary condition is easily adapted to the Lebedev grid and provides an accurate way of truncating the computational domain.

4.2 Perfect Electric Conductors

Perfect electric conductors (PECs) are ideal metals because they are assumed to have infinite conductivity. They are used as a first approximation of metals in circuits especially when low conductive losses are expected. PECs are also used to analytically analyze transmission lines and antennas. Infinite planes of PEC are handled by




Figure 4.2: Reflection coefficient in dB for the two plane wave modes normally incident on a CPML (solid) and measured from a long domain (dashed). The long domain reflection coefficient represents the noise floor of the method for measuring the reflection coefficients given in Appendix B.3. The CPML has not been optimized beyond recommendations in literature but already has sufficient attenuation for the simulations in this thesis.

image theory. That is, the boundary conditions will automatically be satisfied if an appropriate reflection of the sources in the domain are placed on the other side of the PEC. This line of reasoning is the basis for implementing the PEC numerically.

4.2.1 Derivation

At a PEC the following boundary conditions must be satisfied:

$$\frac{\partial \vec{D}}{\partial n} \cdot \hat{n} = 0 \quad , \quad (4.19)$$

$$\vec{E} \times \hat{n} = 0 \quad , \quad (4.20)$$

$$\vec{B} \cdot \hat{n} = 0 \quad , \quad (4.21)$$

$$\frac{\partial \vec{H}}{\partial n} \times \hat{n} = 0 \quad . \quad (4.22)$$

In the case of a positive inward \hat{x} normal PEC at $x = 0$ this becomes

$$\epsilon_{xx} \frac{\partial E_x}{\partial x} + \epsilon_{xy} \frac{\partial E_y}{\partial x} + \epsilon_{xz} \frac{\partial E_z}{\partial x} = 0 \quad , \quad (4.23)$$

$$E_y = 0 \quad , \quad (4.24)$$

$$E_z = 0 \quad , \quad (4.25)$$

$$\mu_{xx} H_x + \mu_{xy} H_y + \mu_{xz} H_z = 0 \quad , \quad (4.26)$$

$$\frac{\partial H_y}{\partial x} = 0 \quad , \quad (4.27)$$

$$\frac{\partial H_z}{\partial x} = 0 \quad . \quad (4.28)$$

All of the components of both the electric and magnetic fields lie on the PEC. Therefore, all of the boundary conditions must be imposed with the same order of accuracy to maintain identical solutions on the four underlying Yee grids.

Consider the electric field gridpoints that lie on the PEC. The tangential components of the electric field are simply zero by (4.24) and (4.25). The normal component is updated using the usual update, (3.12). However, in the case being considered, $\vec{H}[-1, 0, 0]$ lies on the far side of the PEC and may not be accessed. Evaluating (4.27) and (4.28) with central differences gives an expression that reflects the fields across the PEC:

$$H_y[-1, 0, 0] = H_y[1, 0, 0] \quad , \quad (4.29)$$

$$H_z[-1, 0, 0] = H_z[1, 0, 0] \quad . \quad (4.30)$$

Recall that an update in a bulk material only depends on the tangential fields at the neighboring gridpoints. Therefore, (4.29) and (4.30) are sufficient in most cases. However, in certain cases, such as a material interface intersecting a PEC, the normal component must also be reflected. To do this, (4.26) is evaluated with central averages and then combined with (4.29) and (4.30) to give

$$H_x[-1, 0, 0] = -H_x[1, 0, 0] - \frac{2}{\mu_{xx}} (\mu_{xy} H_y[1, 0, 0] + \mu_{xz} H_z[1, 0, 0]) \quad . \quad (4.31)$$

The value of $\vec{H}[-1, 0, 0]$ is then used in the regular update for $E_x[0, 0, 0]$.

Now consider the magnetic field gridpoints. The boundary conditions do not force any of the components to zero. Instead, (4.26) ties the update of the normal component to the tangential components:

$$H_x[0] = \frac{\mu_{xy}}{\mu_{xx}} H_y[0] + \frac{\mu_{xz}}{\mu_{xx}} H_z[0] \quad . \quad (4.32)$$

The tangential components are updated in the same way as the normal electric field.

That is, the regular update is used with the value of the electric field outside the domain replaced using the boundary conditions:

$$E_y[-1, 0, 0] = -E_y[1, 0, 0] \quad , \quad (4.33)$$

$$E_z[-1, 0, 0] = -E_z[1, 0, 0] \quad , \quad (4.34)$$

$$E_x[-1, 0, 0] = E_x[1, 0, 0] - \frac{2}{\epsilon_{xx}} (\epsilon_{xy} E_y[1, 0, 0] + \epsilon_{xz} E_z[1, 0, 0]) \quad . \quad (4.35)$$

The generalization of these formulas to a PEC in the vw plane with pu normal, where $(u, v, w) = (x, y, z)$, (y, z, x) , or (z, x, y) and $p = \pm 1$ are

$$E_v[0] = 0 \quad , \quad (4.36)$$

$$E_w[0] = 0 \quad , \quad (4.37)$$

$$E_u[-p\delta_u] = E_u[p\delta_u] + \frac{2}{\epsilon_{uu}} (\epsilon_{uv} E_v[p\delta_u] + \epsilon_{uw} E_w[p\delta_u]) \quad , \quad (4.38)$$

$$E_v[-p\delta_u] = -E_v[p\delta_u] \quad , \quad (4.39)$$

$$E_w[-p\delta_u] = -E_w[p\delta_u] \quad , \quad (4.40)$$

$$H_u[-p\delta_u] = -H_u[p\delta_u] - \frac{2}{\mu_{uu}} (\mu_{uv} H_v[p\delta_u] + \mu_{uw} H_w[p\delta_u]) \quad , \quad (4.41)$$

$$H_v[-p\delta_u] = H_v[p\delta_u] \quad , \quad (4.42)$$

$$H_w[-p\delta_u] = H_w[p\delta_u] \quad , \quad (4.43)$$

$$H_u[0] = -\frac{1}{\mu_{uu}} (\mu_{uv} H_v[0] + \mu_{uw} H_w[0]) \quad . \quad (4.44)$$

4.2.2 Implementation

The beauty of this derivation is that it replaces unknown field values in an already existing update equation with field values that are already used in the update. In

order to implement a plane of PEC through a point it is sufficient to know the effective material parameters and what the update would be if the PEC did not exist. Knowing this, the implementation can be phrased as a modification to precomputed constants.

Assume that an update is expressed in the form of (3.12) and (3.16), and that \bar{a} and $\bar{\alpha}$ have already been inverted:

$$\vec{E}^{+1}[0] = \left(\bar{a}^{-1}\bar{b}\right) \vec{E}^0[0] + \sum_{p=\pm 1} \sum_{w=x,y,z} (\bar{a}^{-1}\bar{c}[p, w]) \vec{H}^{+\frac{1}{2}}[p\delta_w] \quad , \quad (4.45)$$

$$\vec{H}^{+\frac{1}{2}}[0] = \left(\bar{\alpha}^{-1}\bar{\beta}\right) \vec{H}^{-\frac{1}{2}}[0] + \sum_{p=\pm 1} \sum_{w=x,y,z} (\bar{\alpha}^{-1}\bar{\gamma}[p, w]) \vec{E}^0[p\delta_w] \quad . \quad (4.46)$$

The update equations are a linear combination of the components at the nearest neighboring gridpoints and so are the modifications to the updates (4.36) to (4.44). The terms used to implement the boundary conditions are already in the update equations. Therefore, the inaccessible gridpoints are replaced with their reflections and then the terms are rearranged back into the regular form.

Returning to the case of the PEC with $+x$ normal, the tangential electric field components are maintained at zero by setting

$$\left(\bar{a}^{-1}\bar{b}\right)_{yv} = \delta_{yv} \quad , \quad (\bar{a}^{-1}\bar{c}[p, w])_{yv} = 0 \quad , \quad (4.47)$$

$$\left(\bar{a}^{-1}\bar{b}\right)_{zv} = \delta_{zv} \quad , \quad (\bar{a}^{-1}\bar{c}[p, w])_{zv} = 0 \quad , \quad (4.48)$$

for $v, w = x, y, z$ and $p = \pm 1$. Assuming that the update depends only on the tangential components of the neighboring points, $\vec{H}[-1, 0, 0]$ is reflected into the

domain by setting

$$(\bar{a}^{-1}\bar{c}[+,x])_{uv} = (\bar{a}^{-1}\bar{c}[+,x])_{uv} + (\bar{a}^{-1}\bar{c}[-,x])_{uv} \quad , \quad (4.49)$$

followed by

$$(\bar{a}^{-1}\bar{c}[-,x])_{uv} = 0 \quad , \quad (4.50)$$

for $u = x, y, z$ and $v = yz$. Similarly, $\vec{E}[-1, 0, 0]$ is reflected into the domain by setting

$$(\bar{\alpha}^{-1}\bar{\gamma}[+,x])_{uv} = (\bar{\alpha}^{-1}\bar{\gamma}[+,x])_{uv} - (\bar{\alpha}^{-1}\bar{\gamma}[-,x])_{uv} \quad , \quad (4.51)$$

followed by

$$(\bar{\alpha}^{-1}\bar{\gamma}[-,x])_{uv} = 0 \quad . \quad (4.52)$$

Finally, the boundary condition on the normal magnetic field is implemented by setting

$$(\bar{\alpha}^{-1}\bar{\gamma}[p,w])_{xv} = -\frac{1}{\mu_{xx}} (\bar{\alpha}^{-1}\bar{\gamma}[p,w])_{yv} + (\bar{\alpha}^{-1}\bar{\gamma}[p,w])_{zv} \quad , \quad (4.53)$$

$$(\bar{\alpha}^{-1}\bar{\beta})_{xv} = -\frac{1}{\mu_{xx}} (\bar{\alpha}^{-1}\bar{\beta})_{yv} + (\bar{\alpha}^{-1}\bar{\beta})_{zv} \quad , \quad (4.54)$$

for $v, w = x, y, z$ and $p = \pm 1$.

If the update equations depend on the neighboring normal components then (4.49) becomes

$$(\bar{a}^{-1}\bar{c}[+,x])_{uv} = \begin{cases} (\bar{a}^{-1}\bar{c}[+,x])_{uv} - (\bar{a}^{-1}\bar{c}[-,x])_{uv} & \text{if } v = x, \\ (\bar{a}^{-1}\bar{c}[+,x])_{uv} + (\bar{a}^{-1}\bar{c}[-,x])_{uv} + \frac{2\mu_{xv}}{\mu_{xx}} (\bar{a}^{-1}\bar{c}[-,x])_{ux} & \text{else.} \end{cases} \quad (4.55)$$

Similarly, (4.51) becomes

$$(\bar{\bar{\alpha}}^{-1}\bar{\bar{\gamma}}[+, x])_{uv} = \begin{cases} (\bar{\bar{\alpha}}^{-1}\bar{\bar{\gamma}}[+, x])_{uv} + (\bar{\bar{\alpha}}^{-1}\bar{\bar{\gamma}}[-, x])_{uv} & \text{if } v = x, \\ (\bar{\bar{\alpha}}^{-1}\bar{\bar{\gamma}}[+, x])_{uv} - (\bar{\bar{\alpha}}^{-1}\bar{\bar{\gamma}}[-, x])_{uv} - \frac{2\mu_{xv}}{\mu_{xx}} (\bar{\bar{\alpha}}^{-1}\bar{\bar{\gamma}}[-, x])_{ux} & \text{else.} \end{cases} \quad (4.56)$$

4.2.3 Analysis of a Filled Rectangular Waveguide

The implementation of PEC on a Lebedev grid was tested by simulating a rectangular waveguide. The waveguide was filled with the three materials discussed in Section 2.3. The isotropic filled waveguide was designed to have a cutoff frequency at 4 GHz. The chosen cell dimensions were $\Delta x = \Delta y = \Delta z = h = 0.814$ mm and the number of cells in each direction was $14.5 \times 8.5 \times 1500.5$. These dimensions are the same as the plane wave simulations given in Section 3.9 except that the computational domain is truncated with a CPML absorbing boundary condition discussed in Section 4.1. Each simulation was run at 99% of the calculated maximum timestep. A Yee grid based simulation was also implemented [21, 3] and compared using $\Delta z_{\text{Yee}} = \Delta z_{\text{Leb}}/\sqrt[3]{2}$ to equate the memory usage.

The simulations extracted the wavenumber versus frequency characteristics of the allowed modes using the method described in Appendix B.1. The results (Figure 4.4) show that the simulation reproduces the analytic solution in the isotropic case. There is no analytic solution in the anisotropic case, but the method is certainly stable and physically reasonable. In fact, there are several intuitive features that can be seen in the results. First, the allowed modes converge to two different phase velocities. This makes sense because different modes are analogous to different polarizations and there are two different phase velocities for plane waves in anisotropic media. Furthermore,




Figure 4.3: Setup of computational domain for simulating plane waves, rectangular waveguide, and parallel plate waveguide. The wave type is set by changing the boundary conditions on the x and y normal faces. A pulse is launched down the length of the domain and a line of samples is recorded. The CPML absorbing boundaries prevent unwanted reflections and allow the simulation to be run for a long time.

there are more allowed modes on the diagram because the TE and TM modes are no longer degenerate.

The difference between the Yee grid based results and the Lebedev grid results is less significant than for plane waves. However, The isotropic results suggest that the most significant error is still numerical phase velocity error because the numerical results overestimate the propagation constant. In the anisotropic cases where there is no analytical solution it can only be said that the Lebedev results slightly overshoot the Yee results. Assuming the dominant error is the numerical phase velocity error the Yee grid appears to outperform the Lebedev grid despite the need to extrapolate the values of D on the PEC. As before, this comparison is based solely on memory and accuracy arguments and does not take into account the simple implementation and small stencil used on the Lebedev grid.

Cross-sections of the first two propagating modes (Figure 4.5) were extracted using the method in Appendix B.2. These results are also as expected. The isotropic and




Figure 4.4: Propagation constant as a function of frequency for the modes in a rectangular waveguide simulated using the Lebedev grid (solid) and the Yee grid (dashed). The analytical solution (dotted) is known only for isotropic media. Note the TE_{mn} modes are degenerate with TM_{mn} when $m \neq 0$ and $n \neq 0$. The degeneracy is broken by the anisotropy. Differing asymptotic phase velocities cause mode lines to cross at higher frequencies. All mode lines were obtained from a single simulation run as described in Appendix B.1.

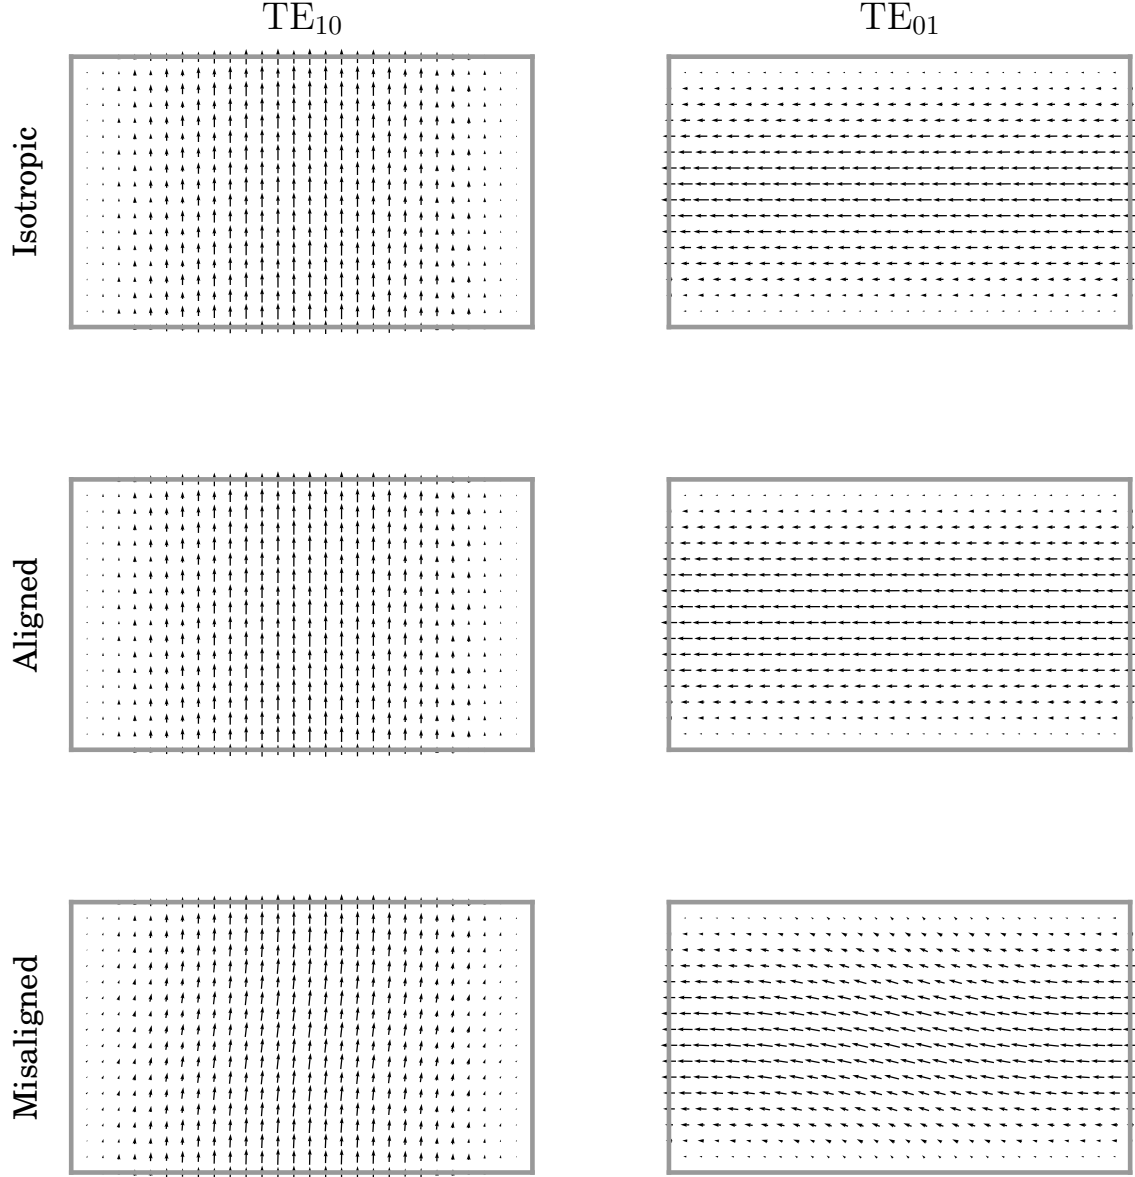


Figure 4.5: Transverse components of the electric field for a plane wave traveling in three different media. The basis polarization states of the uniaxial material do not align with the natural polarizations of the rectangular waveguide in the misaligned case. This conflict causes the fundamental mode to be a TE_{10} perturbed by a TE_{01} , and the second mode to be a TE_{01} perturbed by a TE_{10} . The plots were obtained using the method in Appendix B.2.

aligned simulations both show half sine wave dependence of a single component of the electric field. On the other hand, the first two modes when the filling material was rotated appear as hybrid modes. The first appears as a TE_{10} perturbed by a TE_{01} and the second appears to be the opposite. Again, this result is expected because there is a conflict between the basis linear polarizations in the anisotropic material and the polarizations forced by the boundary conditions of the PEC. This conflict is not observed in the aligned anisotropic case because the two types of polarizations coincide with one another.

4.3 Material Interfaces

Material interfaces are commonly encountered in microwave transmission lines and circuit design. In this thesis, only abrupt discontinuities in the material parameters in a plane the aligns with the computational grid are considered. Such interfaces are commonly encountered in practice, for example, the air-substrate boundary on a circuit board or between layers in an integrated circuit. At the interface between two materials the boundary conditions are well established: (for this section only, the superscript is used to denote which material a field lies within.)

$$\hat{n} \cdot \vec{D}^1 = \hat{n} \cdot \vec{D}^2 \quad , \quad (4.57)$$

$$\hat{n} \times \vec{E}^1 = \hat{n} \times \vec{E}^2 \quad , \quad (4.58)$$

$$\hat{n} \cdot \vec{B}^1 = \hat{n} \cdot \vec{B}^2 \quad , \quad (4.59)$$

$$\hat{n} \times \vec{H}^1 = \hat{n} \times \vec{H}^2 \quad . \quad (4.60)$$

These boundary conditions can be written solely in terms of E and H . For a material interface in the yz plane:

$$\epsilon_{xx}^1 E_x^1 + \epsilon_{xy}^1 E_y^1 + \epsilon_{xz}^1 E_z^1 = \epsilon_{xx}^2 E_x^2 + \epsilon_{xy}^2 E_y^2 + \epsilon_{xz}^2 E_z^2 \quad , \quad (4.61)$$

$$E_y^1 = E_y^2 \quad , \quad (4.62)$$

$$E_z^1 = E_z^2 \quad , \quad (4.63)$$

$$\mu_{xx}^1 H_x^1 + \mu_{xy}^1 H_y^1 + \mu_{xz}^1 H_z^1 = \mu_{xx}^2 H_x^2 + \mu_{xy}^2 H_y^2 + \mu_{xz}^2 H_z^2 \quad , \quad (4.64)$$

$$H_y^1 = H_y^2 \quad , \quad (4.65)$$

$$H_z^1 = H_z^2 \quad . \quad (4.66)$$

4.3.1 Derivation

Assume that there are two sets of gridpoints on the material boundary (Figure 4.6). One set lies just within material 1 and the other just within material 2. These two sets of points are used in the integral derivation of the update equations given in Section 3.7. The boundary conditions (4.61) to (4.66) are then applied so that only the set of points in material 1 need to be stored.

Consider the integral formulation of Maxwell's equations:

$$\oint \vec{H} \cdot d\vec{L} = \iint \left(\bar{\epsilon} \frac{\partial \vec{E}}{\partial t} - \bar{\sigma} \vec{E} + \vec{J} \right) \cdot d\vec{S} \quad , \quad (4.67)$$

$$\oint \vec{E} \cdot d\vec{L} = \iint \left(-\bar{\mu} \frac{\partial \vec{H}}{\partial t} - \bar{\sigma}^* \vec{H} - \vec{M} \right) \cdot d\vec{S} \quad . \quad (4.68)$$

The updates are derived using rectangular integration loops that pass through the nearest neighboring grid points. When considering a point on the boundary, there




Figure 4.6: Grid setup for deriving the update equations at a discontinuous planar material interface. A set of grid points is assumed to lie within each material. After the integral derivation is applied, the boundary conditions are used to relate the fields on each side of the boundary. Therefore, at the interface, only the field values in material 1 need to be stored.

are three orthogonal integration loops. Two loops pass through the material interface and the other lies entirely within material 1. If a section of a loop passes through the material interface then the integration technique is modified to better represent the transition. The line integral is evaluated in two segments. Each segment is approximated by the endpoint rule using the point that lies on the boundary. For example, if the material interface is in the yz -plane at $x = 0$ then a line segment that passes through the interface extends from $-\Delta x/2$ to $\Delta x/2$:

$$\int_{-\frac{\Delta x}{2}}^{\frac{\Delta x}{2}} E_x(x, y, z) dx = \int_{-\frac{\Delta x}{2}}^0 E_x^1(x, y, z) dx + \int_0^{\frac{\Delta x}{2}} E_x^2(x, y, z) dx \quad . \quad (4.69)$$

The line segment in each material is evaluated using the endpoint that lies on the interface:

$$\int_{-\frac{\Delta x}{2}}^{\frac{\Delta x}{2}} E_x(x, y, z) dx \approx \frac{\Delta x}{2} (E_x^1(0, y, z) + E_x^2(0, y, z)) \quad , \quad (4.70)$$

and (4.61) to (4.63) is substituted to remove dependence on the fields in material 2:

$$\int_{-\frac{\Delta x}{2}}^{\frac{\Delta x}{2}} E_x dx \approx \frac{\Delta x}{2} \left(\left(1 + \frac{\epsilon_{xx}^1}{\epsilon_{xx}^2} \right) E_x^1 + \frac{\epsilon_{xy}^1 - \epsilon_{xy}^2}{\epsilon_{xx}^2} E_y^1 + \frac{\epsilon_{xz}^1 - \epsilon_{xz}^2}{\epsilon_{xx}^2} E_z^1 \right) \Big|_{x=0}. \quad (4.71)$$

The approximation for a surface integration follows a similar process:

$$\begin{aligned} \int_{-\frac{\Delta x}{2}}^{\frac{\Delta x}{2}} \int_{-\frac{\Delta y}{2}}^{\frac{\Delta y}{2}} E_z dy dx &= \int_{-\frac{\Delta x}{2}}^0 \int_{-\frac{\Delta y}{2}}^{\frac{\Delta y}{2}} E_z^1(x, y, z) dy dx + \int_0^{\frac{\Delta x}{2}} \int_{-\frac{\Delta y}{2}}^{\frac{\Delta y}{2}} E_z^2(x, y, z) dy dx, \\ &\approx \frac{\Delta x \Delta y}{2} (E_z^1(0, 0, z) + E_z^2(0, 0, z)) \quad , \\ &= \Delta x \Delta y E_z^1(0, 0, z) \quad . \end{aligned} \quad (4.72)$$

Note that there are three surface integrals for each component of the field being updated because $\vec{D} \cdot \hat{n}$ is integrated.

Repeating this for all the required integrals and rearranging the terms gives an expression for the effective material parameters at the interface:

$$\begin{pmatrix} \epsilon_{xx} & \epsilon_{xy} & \epsilon_{xz} \\ \epsilon_{yx} & \epsilon_{yy} & \epsilon_{yz} \\ \epsilon_{zx} & \epsilon_{zy} & \epsilon_{zz} \end{pmatrix} = \begin{pmatrix} \epsilon_{xx}^1 & \epsilon_{xy}^1 & \epsilon_{xz}^1 \\ \frac{\epsilon_{yx}^1}{2} + \frac{\epsilon_{yx}^2}{2} \frac{\epsilon_{xx}^1}{\epsilon_{xx}^2} & \frac{\epsilon_{yy}^1 + \epsilon_{yy}^2}{2} + \frac{\epsilon_{yx}^2}{2} \frac{\epsilon_{xy}^1 - \epsilon_{xy}^2}{\epsilon_{xx}^2} & \frac{\epsilon_{yz}^1 + \epsilon_{yz}^2}{2} + \frac{\epsilon_{yx}^2}{2} \frac{\epsilon_{xz}^1 - \epsilon_{xz}^2}{\epsilon_{xx}^2} \\ \frac{\epsilon_{zx}^1}{2} + \frac{\epsilon_{zx}^2}{2} \frac{\epsilon_{xx}^1}{\epsilon_{xx}^2} & \frac{\epsilon_{zy}^1 + \epsilon_{zy}^2}{2} + \frac{\epsilon_{zx}^2}{2} \frac{\epsilon_{xy}^1 - \epsilon_{xy}^2}{\epsilon_{xx}^2} & \frac{\epsilon_{zz}^1 + \epsilon_{zz}^2}{2} + \frac{\epsilon_{zx}^2}{2} \frac{\epsilon_{xz}^1 - \epsilon_{xz}^2}{\epsilon_{xx}^2} \end{pmatrix}, \quad (4.73)$$

$$\begin{pmatrix} \sigma_{xx} & \sigma_{xy} & \sigma_{xz} \\ \sigma_{yx} & \sigma_{yy} & \sigma_{yz} \\ \sigma_{zx} & \sigma_{zy} & \sigma_{zz} \end{pmatrix} = \begin{pmatrix} \sigma_{xx}^1 & \sigma_{xy}^1 & \sigma_{xz}^1 \\ \frac{\sigma_{yx}^1}{2} + \frac{\sigma_{yx}^2}{2} \frac{\epsilon_{xx}^1}{\epsilon_{xx}^2} & \frac{\sigma_{yy}^1 + \sigma_{yy}^2}{2} + \frac{\sigma_{yx}^2}{2} \frac{\epsilon_{xy}^1 - \epsilon_{xy}^2}{\epsilon_{xx}^2} & \frac{\sigma_{yz}^1 + \sigma_{yz}^2}{2} + \frac{\sigma_{yx}^2}{2} \frac{\epsilon_{xz}^1 - \epsilon_{xz}^2}{\epsilon_{xx}^2} \\ \frac{\sigma_{zx}^1}{2} + \frac{\sigma_{zx}^2}{2} \frac{\epsilon_{xx}^1}{\epsilon_{xx}^2} & \frac{\sigma_{zy}^1 + \sigma_{zy}^2}{2} + \frac{\sigma_{zx}^2}{2} \frac{\epsilon_{xy}^1 - \epsilon_{xy}^2}{\epsilon_{xx}^2} & \frac{\sigma_{zz}^1 + \sigma_{zz}^2}{2} + \frac{\sigma_{zx}^2}{2} \frac{\epsilon_{xz}^1 - \epsilon_{xz}^2}{\epsilon_{xx}^2} \end{pmatrix}. \quad (4.74)$$

Terms of the form seen in (4.70) change some of the curl matrices. Note that the material interface boundary condition applied at the neighboring points on the

interface plane couples all of the field components together. For example,

$$\bar{c}[+, y] = \begin{pmatrix} 0 & 0 & \frac{1}{\Delta y} \\ 0 & 0 & 0 \\ -\frac{1}{2\Delta y} \left(1 + \frac{\mu_{xx}^1}{\mu_{xx}^2}\right) & -\frac{1}{2\Delta y} \frac{\mu_{xy}^1 - \mu_{xy}^2}{\mu_{xx}^2} & -\frac{1}{2\Delta y} \frac{\mu_{xz}^1 - \mu_{xz}^2}{\mu_{xx}^2} \end{pmatrix}. \quad (4.75)$$

Therefore, if one of the materials has off-diagonal elements in the tensor then the updates will depend on the normal component of the neighboring gridpoints.

4.3.2 Implementation

A material interface is implemented in software by computing \bar{a} , \bar{b} , $\bar{\alpha}$, and $\bar{\beta}$ from the effective material parameters, modifying the curl matrices, and then inverting \bar{a} and $\bar{\alpha}$. The difficulty is that \bar{a} and $\bar{\alpha}$ cannot be inverted until after the curl matrices have been modified. It follows that the modifications cannot be expressed as generally as they were for PEC and the intersection of three or more materials must be handled as a special case. These derivations are left as future work.

The equations derived above can be generalized to a boundary with an $\hat{s} \in \{\hat{x}, \hat{y}, \hat{z}\}$ normal:

$$\sigma_{uv} = \begin{cases} \sigma_{uv}^1 & u = s \\ \frac{\sigma_{uv}^1}{2} + \frac{\sigma_{uv}^2}{2} \frac{\epsilon_{ss}^1}{\epsilon_{ss}^2} & u \neq s, v = s \\ \frac{\sigma_{uv}^1 + \sigma_{uv}^2}{2} + \frac{\sigma_{us}^2}{2} \frac{\epsilon_{sv}^1 - \epsilon_{sv}^2}{\epsilon_{ss}^2} & u \neq s, v \neq s \end{cases}. \quad (4.76)$$

The effective permittivity is found by replacing σ_{uv} with ϵ_{uv} in (4.76). Similarly, the effective magnetic material parameters are found by replacing σ_{uv} by σ_{uv}^* and ϵ_{uv} by

μ_{uv} . At a material interface, the curl matrices are given by

$$c_{uv}[p, w] = \begin{cases} \varepsilon_{uvw} \frac{p}{2\Delta w} \left(1 + \frac{\mu_{ss}^1}{\mu_{ss}^2} \right) & u, w \neq s, v = s \\ \varepsilon_{uvw} \frac{p}{2\Delta w} \left(1 + \frac{\mu_{sv}^1 - \mu_{sv}^2}{\mu_{ss}^2} \right) & u, v, w \neq s \\ \varepsilon_{uvw} \frac{p}{\Delta w} & w = s \text{ or } u = s \end{cases}. \quad (4.77)$$

The expression for $\gamma_{uv}[p, w]$ is found by replacing μ_{uv} with ϵ_{uv} .

On the Yee grid, a planar material interface is handled by averaging the material parameters at the location of the interface [22]. The interpolation to enforce the constitutive relations then uses material parameters from surrounding locations in order to ensure stability [21]. Therefore, the fields in each material are indistinguishable and three planes of gridpoints need modified updates to enforce the boundary conditions. In contrast, the effective material parameters on the Lebedev grid coincide with the modifications to the curl matrices such that the field value on both sides of the interface is well defined. Furthermore, only gridpoints lying in the plane of the interface are affected.

4.3.3 Reflection and Transmission of Normally Incident Plane Waves

The material interface was tested by measuring the reflection and transmission coefficients of plane waves normally incident on a slab of isotropic dielectric material embedded in the three materials introduced in Section 2.3. The simulated dielectric slab is thin and thus the two material interfaces are relatively close together. An incident pulse reflects between the two interfaces multiple times and the pulses interfere with one another. Intuitively, the reflection and transmission coefficients are a function of the ratio between the wavelength and the thickness of the slab. Therefore,

this simulation proves that the material boundary conditions hold for all simulated frequencies and that the boundary lies at the expected location.

The reflection coefficient at a planar anisotropic interface cannot truly be found analytically. However, it is straightforward to setup a system of equations relating the reflection and transmission coefficients of the allowed polarizations by forcing the continuity of the tangential components of \vec{E} and \vec{H} . Assume that the incident polarization is an eigensolution in the incident medium. The transmitted wave will be a linear combination of the eigensolutions in the second material and the reflection will be a linear combination of the eigensolutions in the incident medium. Let r_1 , r_2 be the reflection coefficient of the two polarizations in the incident medium and t_1 , t_2 be the transmission coefficient of the two polarizations in the transmitted medium. Also, let E_x^{r1} be the x -component of the electric field for the first polarization of a plane wave traveling away from the interface. With similar definitions for the other fields, continuity of the tangential components requires

$$E_x^i + r_1 E_x^{r1} + r_2 E_x^{r2} = t_1 E_x^{t1} + t_2 E_x^{t2} \quad , \quad (4.78)$$

$$E_y^i + r_1 E_y^{r1} + r_2 E_y^{r2} = t_1 E_y^{t1} + t_2 E_y^{t2} \quad , \quad (4.79)$$

$$H_x^i + r_1 H_x^{r1} + r_2 H_x^{r2} = t_1 H_x^{t1} + t_2 H_x^{t2} \quad , \quad (4.80)$$

$$H_y^i + r_1 H_y^{r1} + r_2 H_y^{r2} = t_1 H_y^{t1} + t_2 H_y^{t2} \quad . \quad (4.81)$$

Assuming a normally incident plane wave, all of the components of electric and magnetic fields can be found from the material properties and the known direction of propagation as shown in Section 2.2. Equations (4.78) to (4.81) are then rearranged into a matrix and the reflection and transmission coefficients are found by invert-




Figure 4.7: Setup of the computational domain for measuring the reflection and transmission coefficients of a dielectric interface. A pulse is launched down the length of the domain and a line of samples is recorded before and after the dielectric slab. The two line samples are transformed along z and t so that the reflection and transmission coefficients of each mode may be found as a function of frequency as described in Appendix B.3.

ing the matrix. If the second material is isotropic then the polarization of the two orthogonal transmitted states is arbitrary and one can be matched to the incident polarization. Matching the polarizations of \vec{E}^i , \vec{E}^{r1} and \vec{E}^{t1} guarantees that $r_2 = t_2 = 0$. If this is not enforced then a simulation would be required for each incident polarization.

Since the two polarizations are not coupled by reflection at a single interface, the total reflection (Γ) and transmission (T) coefficients are also uncoupled. Therefore, the analytical solution for each polarization takes the same form as the analytical solution for an isotropic slab in an isotropic background medium. The total reflection coefficient is determined by the reflection coefficient at the first and second slab interface (r_{1a} , r_{1b}), the transmission coefficient at the first and second slab interface (t_{1a} , t_{1b}), the propagation constant within the slab (k_{t1}), and the slab thickness (d)

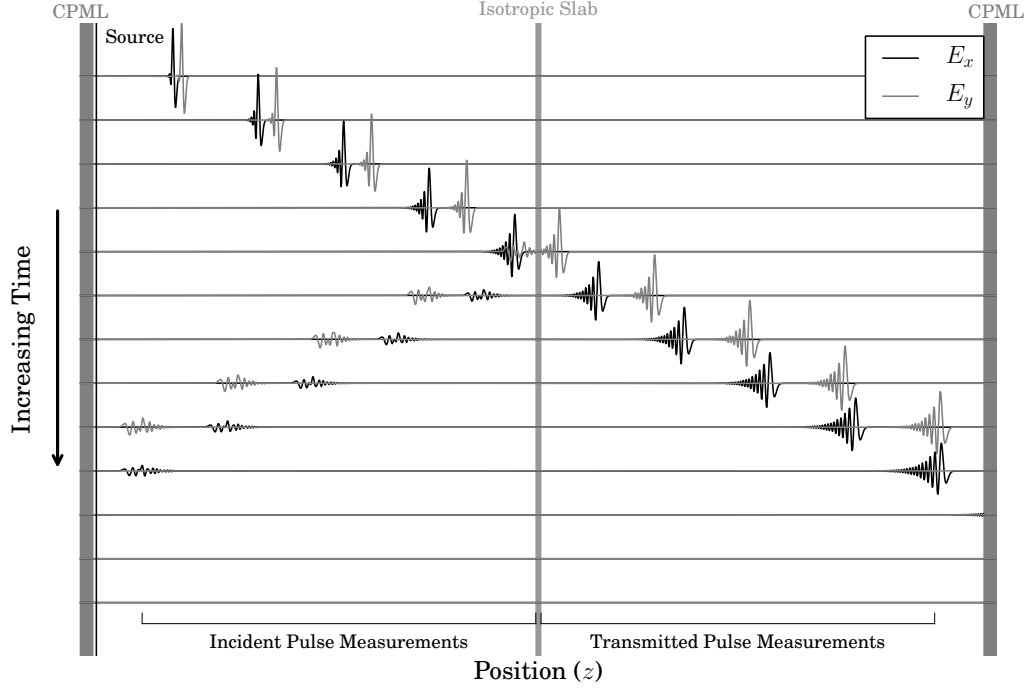


Figure 4.8: Pulse evolution over time for a plane wave normally incident on a slab of isotropic material embedded in sapphire with a permittivity tensor aligned with the grid. A line of samples is recorded in the sapphire on either side of the slab. Both lines are transformed to frequency and wavenumber to measure reflection and transmission (Figure 4.9).

[17]:

$$\Gamma = \frac{r_{1a} + r_{1b}e^{j2k_{t1}d}}{1 + r_{1a}r_{1b}e^{j2k_{t1}d}} \quad (4.82)$$

$$T = \frac{t_{1a}t_{1b}e^{j2k_{t1}d}}{1 + r_{1a}r_{1b}e^{j2k_{t1}d}} \quad (4.83)$$

The numerical simulation was setup as shown in Figure 4.7 with the simulation parameters the same as Section 4.2.3. The reflection and transmission coefficients were calculated as a function of propagating mode and frequency using the method in Appendix B.3. The results are plotted in Figure 4.10. The Yee grid method was




Figure 4.9: Electric field amplitude in frequency and wavenumber for plane waves normally incident on a dielectric slab. The spatial Fourier transform isolates the forwards and backwards propagating waves. The reflection and transmission coefficients are measured by zooming into each polarization at a fixed frequency using the chirp Z-transform and comparing the amplitude of the incident, reflected, and transmitted waves.

based on the interpolation algorithm proposed by Werner [21] with the improved material averaging scheme proposed by Oskooi [22].

There are two errors visible in the simulated reflection coefficients in Figure 4.10. The first error is in the location of the transmission maxima and can be attributed to numerical phase velocity error. This error is less significant on the Yee grid because it is better sampled. However, there is a second error of comparable magnitude in the Yee grid results but not in the Lebedev results. Notice that the Yee grid results underestimate the reflection coefficient at high frequencies. This occurs because the




Figure 4.10: Magnitude of the reflection (Γ) and transmission (T) coefficients for an isotropic slab ($\epsilon = 22, d = 11.8$ mm) embedded in the three media described in Section 2.3 simulated using the Lebedev grid (solid) and the Yee grid (dashed). There are two visible errors between the numerical and analytical curves (dotted): Numerical phase velocity error (downshift in frequency) and material averaging error (underestimated reflection coefficient). The Lebedev grid does not underestimate $|\Gamma|$ but has greater numerical dispersion.

reflectivity of each dielectric interface is decreased by the averaged material parameters. This effect is most noticeable when the wavelength becomes small compared to the cell size. Therefore, the Lebedev grid produces results of comparable accuracy when compared to the Yee grid using equal memory arguments. Additionally, the Lebedev grid is simpler to implement and requires less calls to memory.

4.4 Planar Boundary Conditions Summary

In this chapter, planar PEC boundaries, planar material interfaces, and CPML absorbing boundary conditions were implemented on the Lebedev grid. Infinite planes of perfect electric conductor (PEC) were implemented using the principle of image theory. The reflection was slightly modified to take into account the anisotropy and only store \vec{E} and \vec{H} . The equations were rearranged into a modification of a pre-computed update equation so that the intersection of multiple PECs and/or material interfaces can be implemented in the same way as a regular PEC-material boundary. The PEC implementation was verified by extracting the wavenumber versus frequency characteristics of the allowed modes in a rectangular waveguide.

Update equations at a planar material interface were also derived. The boundary conditions were rearranged into effective material parameters and a modification to the curl matrices. The FDTD material interface was verified by measuring the reflection and transmission coefficients of plane waves incident on an isotropic dielectric slab embedded in an infinite anisotropic medium. Unlike other methods [3, 4], the implementation of both PEC and material interfaces can be written solely in terms of the \vec{E} and \vec{H} fields that are accessed in the regular update equations. Furthermore, the implementations are stable to the same Courant limit as the effective material at

the boundary location.

Finally, the convolutional perfectly matched layer (CPML) absorbing boundary condition was introduced as a means of truncating the computational domain. The formulation on the Lebedev grid is the same as for the Yee grid and anisotropy is easily taken into account because the CPML formulation is material independent. The CPML reflection coefficient was measured and compared to the noise floor of the method for extracting reflection coefficients. At low frequencies the CPML attenuates down to the noise floor but the reflection coefficient rises at high frequencies. Regardless, the reflected power is less than -40 dB (0.01%) for the entire frequency range of interest and is sufficient for the simulations in this thesis.

At this point the simulation tool has been developed far enough to be of practical use for transmission line and antenna design. From here on, any additional components that need to be simulated are application dependent and their implementation is left as future work. Instead of further developing the Lebedev grid, the next chapter shows how the principles of the Lebedev grid can be used to discretize other sets of differential equations.

Chapter 5

Grid Extensions

The Lebedev grid can be viewed as the central difference based extension to the Yee grid for anisotropic materials. That is, the grid structure is chosen to be the most sparse discretization scheme that uses only central differences to approximate the governing equations. This chapter applies this concept to solve other sets of differential equations that describe material behavior in electromagnetics. The first example describes the behavior of ferrites by simultaneously solving Gilbert's equation of motion. The second example reduces the required memory of the Lebedev grid and applies in the special case of a dielectric material with symmetry about the z -axis in the computational domain.

5.1 Ferrites Simulated by Gilbert's Equation of Motion

Ferrites and magnetized plasmas are both dispersive gyrotropic materials that are best simulated by taking into account an auxiliary differential equation in addition to Maxwell's curl equations [24, 30, 31, 32]. The equations being considered are Maxwell's curl equations with anisotropic permittivity ($\bar{\epsilon}$), conductivity ($\bar{\sigma}$), and

scalar permeability (μ):

$$\nabla \times \vec{h} = \bar{\bar{\epsilon}} \frac{\partial \vec{E}}{\partial t} + \bar{\sigma} \vec{E} \quad , \quad (5.1)$$

$$\nabla \times \vec{E} = -\mu \frac{\partial \vec{h}}{\partial t} - \mu \frac{\partial \vec{m}}{\partial t} \quad , \quad (5.2)$$

and the lossless Gilbert's equation for small signals with the saturation magnetization (\vec{M}_s) and biasing magnetic field (\vec{H}_0) oriented in arbitrary directions [24]:

$$\frac{\partial \vec{m}}{\partial t} = -\gamma \left(\mu \vec{M}_s \times \vec{h} + \vec{m} \times \mu \vec{H} \right) \quad , \quad (5.3)$$

where γ is the gyromagnetic ratio of an electron. Here it is assumed that the magnetic field, $\vec{H} = \vec{H}_0 + \vec{h}$, and the magnetization, $\vec{M} = \vec{M}_s + \vec{m}$, Can be divided into a time invariant, curl free vector (\vec{H}_0 , \vec{M}_s) and a small signal perturbation (\vec{h} , \vec{m}).

Notice that the time derivatives of \vec{h} and \vec{m} are coupled to each other but are independent of \vec{E} . This suggests that (5.2) and (5.3) would be easier to discretize if they were written as a single equation. As in Chapter 3, the discretized curl and cross product will be represented by a 3 by 3 matrix to vector multiplication, $\bar{\bar{I}}$ is a 3 by 3 identity matrix and $\bar{\bar{0}}$ is a 3 by 3 zero matrix:

$$\begin{pmatrix} -\mu \bar{\bar{I}} & -\mu \bar{\bar{I}} \\ \bar{\bar{0}} & \bar{\bar{I}} \end{pmatrix} \frac{\partial}{\partial t} \begin{pmatrix} \vec{h} \\ \vec{m} \end{pmatrix} = \begin{pmatrix} \nabla \times \\ \bar{\bar{0}} \end{pmatrix} \begin{pmatrix} \vec{E} \end{pmatrix} + \begin{pmatrix} \bar{\bar{0}} & \bar{\bar{0}} \\ -\gamma \mu \vec{M}_s \times & \gamma \mu \vec{H}_0 \times \end{pmatrix} \begin{pmatrix} \vec{h} \\ \vec{m} \end{pmatrix} \quad . \quad (5.4)$$

Let \vec{r} be a 6D vector containing both \vec{h} and \vec{m} . Then (5.4) can be rewritten as

$$\bar{\bar{A}}^* \frac{\partial \vec{r}}{\partial t} = \bar{\bar{B}}^* \vec{r} + \sum_{w=xyz} \sum_{p=\pm 1} \bar{\bar{C}}^* \vec{E} \quad , \quad (5.5)$$




Figure 5.1: Lebedev (left) and Yee (right) grids for simulating magnetized ferrites. For both grids the magnetization vector (\vec{m}) and small signal magnetic field vector (\vec{h}) are collocated. Coupling between orthogonal components of \vec{m} and \vec{h} is easier to handle on the Lebedev grid because all field components are collocated.

where \bar{A}^* and \bar{B}^* are 6 by 6 matrices and \bar{C}^* is a 3 by 6 matrix of spatial derivatives that represents the curl.

Equations (5.1) and (5.5) are then discretized using central differences and averages:

$$\left(\frac{\bar{\epsilon}}{\Delta t} + \frac{\bar{\sigma}}{2}\right) \vec{E}^{+1}[0] = \left(\frac{\bar{\epsilon}}{\Delta t} - \frac{\bar{\sigma}}{2}\right) \vec{E}^0[0] + \sum_{p=\pm 1} \sum_{w=xyz} \bar{c}[p, w] \vec{h}^{+\frac{1}{2}}[p\delta_w] \quad , \quad (5.6)$$

$$\left(\frac{\bar{A}^*}{\Delta t} + \frac{\bar{B}^*}{2}\right) \vec{r}^{+\frac{1}{2}}[0] = \left(\frac{\bar{A}^*}{\Delta t} - \frac{\bar{B}^*}{2}\right) \vec{r}^{-\frac{1}{2}}[0] + \sum_{p=\pm 1} \sum_{w=xyz} \bar{\gamma}[p, w] \vec{E}^0[p\delta_w] \quad . \quad (5.7)$$

Here $\bar{c}[p, w]$ and $\bar{\gamma}[p, w]$ are the discrete representation of the curl:

$$c_{uv}[p, w] = \gamma_{uv}[p, w] = \epsilon_{uvw} \frac{p}{\Delta w} \quad , \quad (5.8)$$

except that $\bar{\gamma}[p, w]$ is a 6 by 3 matrix with the upper 3 by 3 block given by (5.8) and

the lower block is a zero matrix. The matrices on the left hand side of (5.6) and (5.7) depend on the biasing fields and material parameters and can be inverted to give a recursive updating scheme for \vec{E} , \vec{h} and \vec{m} .

The update equations, (5.6) and (5.7), relate all of the components of \vec{E} to all the components of the neighboring \vec{h} and \vec{m} fields in the positive and negative x , y , and z -directions and vice versa. It follows that the grid should have \vec{E} staggered from both \vec{h} and \vec{m} in time and space but all of the components should be stored at each location. All of the components of \vec{h} and \vec{m} need to be collocated, and staggered from \vec{E} in the cardinal directions. The resulting grid shown in Figure 5.1 is the same as that proposed in Chapter 3 with \vec{m} stored wherever \vec{h} is stored.

As seen for anisotropic materials, the Lebedev based grid can be decomposed into several Yee based grids that propagate uncoupled solutions when all of the materials are isotropic and the biasing field aligns with one of the coordinate axes. Therefore, distributed grid excitations and averaged sampling must also be used on this grid. Similarly, the benefits observed on the Lebedev grid are maintained. For example, this approach uses a localized stencil and provides sufficient information about the field values to accurately enforce boundary conditions.

5.2 Anisotropic Dielectrics With an Axis of Symmetry

A second example deals with the memory limitation of the Lebedev grid by considering a special case of dielectric anisotropy. When the dielectric is symmetric about a cardinal axis then it is sufficient to use a grid that is a superposition of two shifted Yee grids instead of four as shown in Figure 5.2. This grid still approximates all of the spatial derivatives that are coupled in by the off-diagonal entries in the permittivity




Figure 5.2: Reduced Lebedev (right) and Yee (left) grids for simulating anisotropic dielectrics with symmetry about the z -axis. In this special case it is sufficient to represent the solution using a superposition of two Yee grids. This reduces the memory requirements of the Lebedev grid and still allows Maxwell's equations with constitutive relations to be evaluated with central differences.

tensor with central derivatives, but does so with half as much memory.

Consider a permittivity tensor of the form

$$\bar{\epsilon} = \begin{pmatrix} \epsilon_{xx} & \epsilon_{xy} & 0 \\ \epsilon_{yx} & \epsilon_{yy} & 0 \\ 0 & 0 & \epsilon_{zz} \end{pmatrix} . \quad (5.9)$$

An example of a permittivity tensor with this form is an anisotropic dielectric with its principle axis frame rotated about a single cardinal axis in the computational domain. In this case only E_x and E_y are related by the anisotropy and need to be collocated. E_z can be staggered from the other two components. Every component needs to be part of a Yee grid in order to be updated by the curl equations. Therefore, this set of differential equations can be solved with two Yee grids shifted such that E_x and E_y are collocated.

This reduced grid has some key differences and similarities to both the Lebedev and Yee grids. The Yee grid does not have collocated field components and significant errors arise when physical situations couple orthogonal field components together, for example, at material discontinuities. The reduced Lebedev grid is halfway in between. Some fields are collocated and therefore a greater number of physical situations will be accurately taken into account than the Yee grid but less than can be handled by the Lebedev grid. Conversely, the reduced grid uses twice the memory of the regular Yee grid and therefore has worse numerical dispersion than the Yee grid but better dispersion error than the Lebedev grid.

5.3 Grid Extensions Summary

In this chapter two grid structures have been proposed for other types of material parameters. The analysis shows that central difference based approaches trade off higher memory requirements and complicated grid excitation for small concise stencils and robust stability criteria. The standard approach to analyzing materials described by complex constitutive relations is to choose a grid discretization and then evaluate all required differential operators on that grid. The central difference based approach is to first approximate the set of differential equations with finite differences and then work backwards to uncover the most sparse grid that stores all the required field values. Experience suggests that this grid will likely have degeneracies when the complexity of the constitutive relations is relaxed. The key to understanding this degeneracy lies in the integral derivation of Section 3.7.

Chapter 6

Conclusions

This thesis illustrated the development of an FDTD simulation tool based on central differences for anisotropic media by changing the grid discretization from the Yee grid to the Lebedev grid. Chapter 2 introduced anisotropic media and the FDTD method. Analytical techniques and modern Yee grid based methods were presented and used for comparison throughout the development of the Lebedev grid. Chapter 3 covered the derivation of the updating algorithm and an analysis of the computational errors. That is, a dispersion relation was found with subsequent proof of the stability criterion, conservation of charge on the discrete grid was proven, and the degeneracy of the Lebedev grid was dealt with by appropriate grid excitation. Chapter 4 added planes of PEC, planar discontinuous material interfaces, and CPML absorbing boundary conditions. All of these were implemented using the same compact stencil as the update equations in an infinite domain (excepting the CPML auxiliary functions). Chapter 5 showed how the concept of the Lebedev grid can be extended to other materials in electromagnetism using a central difference based approach. In particular, the grid structure for simulating magnetized ferrites using Gilbert's equation of motion and the grid structure for the special case of an anisotropic dielectric with symmetry about the z -axis was given.

Numerical simulations extracted the z -component of the propagation constant

over frequency for plane waves launched from a periodic array of dipoles in an anisotropic dielectric and for a rectangular waveguide filled with an anisotropic dielectric. The reflection and transmission coefficients over frequency of an isotropic slab embedded in an anisotropic dielectric was also given. A Yee grid scheme [21, 3, 22] reproduced the Lebedev results for all simulations. The Yee grid used a smaller cell size so that both grids discretized the same 3D volume with the same amount of memory. It was found that for the dipole array and rectangular waveguide simulations the numerical phase velocity error was the largest computational error. This error is smaller on the Yee grid because a smaller cell size was used. However, the Lebedev grid produced more accurate results for the dielectric slab simulation because the Yee grid has difficulty representing dielectric discontinuities [21, 22, 23]. Therefore, the phase velocity error is not the most significant computational error and the Lebedev grid achieves better accuracy with a larger cell size. This justifies the increased memory cost. In summary, it was found that the Yee grid has lower numerical dispersion but the Lebedev grid can better represent discontinuities in the material parameters. Additionally, the Lebedev grid has a small stencil that does not change at material interfaces or metal boundaries and therefore has a smaller compute time and is easier to implement.

6.1 Accomplishments

The work in this thesis is related to several publications. The work began with analyzing dispersion properties of the FDTD method on a face-centered cubic grid for the scalar wave equation. This led to a conference presentation at the Antennas and Propagation Society 2011 meeting [33] and subsequent paper in the Journal of

Computational Physics[34]. That work was extended to Maxwell's equations and has been submitted to Transactions on Antennas and Propagation [35]. The Lebedev grid work flows from the FCC grid work because the electric and magnetic fields each form their own FCC grid. However, the stencil on the Lebedev grid is more closely related to the Yee grid stencil and this led to the notion of a central difference based approximation to a set of differential equations. The Lebedev grid work has been presented at the 2012 Antennas and Propagation Society meeting [36] and has been submitted to Transactions on Antennas and Propagation[37]. Finally, the application of the central difference based approximation of the ferrite differential equations has been presented at the International Conference on Electromagnetics in Advanced Applications[38].

This thesis achieved the goals outlined in Chapter 1 as follows:

- The structure of the Lebedev grid was found by discretizing Maxwell's equations with tensor constitutive relations using only central differences. A similar derivation yielded the required grid structure for solving Gilbert's equation of motion and for dielectric tensors that have symmetry about the z -axis.
- A dispersion relation was derived for the Lebedev grid and the maximum stable timestep was found from an eigenvalue equation involving the constitutive relations and the cell dimensions.
- It was proven that the Lebedev grid has a consistent discrete calculus and thus charge is conserved on the grid.
- It was shown that the Lebedev grid supports spurious solutions and these unphysical solutions are related to the decomposition of the Lebedev grid into four Yee grids. It was further shown that these spurious solutions are eliminated by

using a distributed excitation that excites all four Yee grids equivalently.

- It was found that the numerical phase velocity error was the most significant error for all example simulations. In contrast, the Yee grid exhibits material averaging errors at discontinuous material interfaces.
- Planes of perfect electric conductor were introduced using the same form of update equations as used in a bulk material. In contrast, the stencil on the Yee grid reaches two grid cells away from the plane.
- The update on the interface at a material discontinuity was derived and shown to be in the same form as the update equations in a bulk material. The field values on both sides of the interface are known using the boundary conditions. The Lebedev method was shown to better represent material discontinuities, even in the isotropic case.
- A CPML absorbing boundary condition was implemented with the same computational complexity and accuracy as a Yee grid based CPML.

These accomplishments show that the Lebedev grid is better suited than the Yee grid to problems that are limited by an error other than the numerical phase velocity error. It follows that the central difference based approach to solving other sets of differential equations is worth considering. Additionally, further investigation into the Lebedev grid is required to elucidate more situations where the Lebedev grid outperforms the Yee grid. These situations likely take advantage of the collocated field components on the Lebedev grid. Specific suggestions are given in the next section.

6.2 Future Work

- It was found that numerical dispersion error dominates Lebedev grid simulations. Adapting methods of reducing dispersion, such as [39, 40], to the Lebedev grid would significantly increase its applicability.
- Equations for planar PEC have been presented. These planes can only be used in the bulk of the domain (for example, as microstrip lines) once the formulas for the edge singularity are derived. Note that a normal component will need to be stored on either side of the plane and will require modification of the grid storage structure.
- Similarly, only planar material discontinuities have been considered. Considering material interfaces that do not align with the computational grid, materials with graded parameters, material edges, and the intersection of multiple materials are all possible avenues of research.
- Continued development of the two grids presented in Chapter 5 as well as central difference based approaches to other sets of differential equations (For example, for magnetized plasma) will likely show improvements over Yee grid based methods.
- A final avenue of research would be to use the Lebedev grid as a subgrid to the Yee grid. The crucial requirement at the interface would be relating all of the degenerate Yee grids on the Lebedev grid to the single Yee grid. A possible configuration to help this issue would be to shift the single Yee grid by a quarter cell in each direction so that it is equidistant from all degenerate grids.

Bibliography

- [1] K. S. Yee, “Numerical solution of initial boundary value problems involving Maxwell’s equations in isotropic media,” *IEEE Trans. Antennas Propag.*, vol. 14, pp. 302–307, 1966.
- [2] S. G. Garcia, T. M. Hung-Bao, R. G. Martin, and B. G. Olmedo, “On the application of finite methods in time domain to anisotropic dielectric waveguides,” *IEEE Trans. Microw. Theory Tech.*, vol. 44, no. 12, pp. 2195–2206, 1996.
- [3] A. P. Zhao, J. Juntunen, and A. Raisanen, “An efficient FDTD algorithm for the analysis of microstrip patch antennas printed on a general anisotropic dielectric substrate,” *IEEE Trans. Microw. Theory Tech.*, vol. 47, no. 7, pp. 1142–1146, 1999.
- [4] C. Kung and R. Lee, “Alternative FDTD updating schemes for anisotropic materials,” *IEEE Trans. Antennas Propag.*, vol. 57, no. 12, pp. 3903–3910, 2009.
- [5] S. Davydycheva, V. Druskin, and T. Habashy, “An efficient finite-difference scheme for electromagnetic logging of 3D anisotropic inhomogeneous media,” *Geophysics*, vol. 68, no. 5, pp. 1525–1535, 2003.
- [6] V. Lisitsa and D. Vishnevskiy, “Lebedev scheme for the numerical simulation of wave propagation in 3D anisotropic elasticity,” *Geophysical Prospecting*, vol. 58,

- no. 4, pp. 619–635, 2010.
- [7] H. Bernth and C. Chapman, “A comparison of the dispersion relations for anisotropic elastodynamic finite-difference grids,” *Geophys.*, vol. 76, no. 3, pp. 43–49, 1011.
 - [8] V.I. and Lebedev, “Difference analogues of orthogonal decompositions, basic differential operators and some boundary problems of mathematical physics. i,” *Soviet Comput. Maths. Math. Phys.*, vol. 4, no. 3, pp. 69–92, 1964.
 - [9] —, “Difference analogues of orthogonal decompositions, basic differential operators and some boundary problems of mathematical physics. ii,” *Soviet Comput. Maths. Math. Phys.*, vol. 4, no. 4, pp. 36–50, 1964.
 - [10] D. Lioubtchenko, S. Dudorov, J. Mallat, J. Tuovinen, and A. Räisänen, “Low-loss sapphire waveguides for 75–110 GHz frequency range,” *IEEE Microw. Wireless Compon. Lett.*, vol. 11, no. 6, pp. 252–254, 2001.
 - [11] J. C. Rautio, “A proposed uniaxial anisotropic dielectric measurement technique,” *IEEE MTT-S Workshop on Signal Integrity and High-Speed Interconnects*, pp. 59–62, 2009.
 - [12] C. G. Someda, *Electromagnetic Waves*, 2nd ed. Boca Raton, FL: CRC Press, 2006.
 - [13] A. Taflove and S. C. Hagness, *Computational Electromagnetics: The Finite-Difference Time-Domain*, 3rd ed. Boston, MA: Artech House, 2005.
 - [14] R. A. Chilton and R. Lee, “The lobatto cell: Robust, explicit, higher order FDTD that handles inhomogeneous media,” *IEEE Trans. Antennas Propag.*, vol. 56, no. 8, pp. 2167–2177, 2008.
 - [15] R. E. Collin, *Field Theory of Guided Waves*, 2nd ed. Piscataway, NJ: Oxford

University Press, 1990.

- [16] ———, *Foundations of Microwave Engineering*, 2nd ed. Hoboken, NJ: John Wiley and Sons, 1992.
- [17] W. Chew, *Waves and Fields in Inhomogeneous Media*. Hoboken, NJ: Wiley, 1999.
- [18] D. R. Lide, Ed., *CRC Handbook of Chemistry and Physics*, 88th ed. Boca Raton, FL: CRC Press, 2008.
- [19] F. L. Teixeira, “Time-domain finite-difference and finite-element methods for Maxwell equations in complex media,” *IEEE Trans. Antennas Propag.*, vol. 56, no. 8, pp. 2150–2166, 2008.
- [20] J. Schneider and S. Hudson, “A finite-difference time-domain method applied to anisotropic material,” *IEEE Trans. Antennas Propag.*, vol. 41, no. 7, pp. 994–999, 1993.
- [21] G. R. Werner and J. R. Cary, “A stable FDTD algorithm for non-diagonal, anisotropic dielectrics,” *J. Comp. Phys.*, vol. 226, no. 1, pp. 1085–1101, 2007.
- [22] A. F. Oskooi, C. Kottke, and S. G. Johnson, “Accurate finite-difference time-domain simulation of anisotropic media by subpixel smoothing,” *Optics Letters*, vol. 34, no. 18, pp. 2778–2780, 2009.
- [23] C. A. Bauer, G. R. Werner, and J. R. Cary, “A second-order 3D electromagnetics algorithm for curved interfaces between anisotropic dielectrics on a yee mesh,” *J. Comp. Phys.*, vol. 230, no. 5, pp. 2060–2075, 2011.
- [24] M. Okoniewski and E. Okoniewska, “FDTD analysis of magnetized ferrites: a more efficient algorithm,” *IEEE Microw. Guided Wave Lett.*, vol. 4, no. 6, pp. 169–171, 1994.

- [25] V. Lisitsa and D. Vishnevskiy, “On specific features of the Lebedev scheme in simulating elastic wave propagation in anisotropic media,” *Numerical Analysis and Applications*, vol. 4, no. 2, pp. 125–135, 2011.
- [26] J. M. Hyman and M. Shashkov, “Natural discretizations for the divergence, gradient, and curl on logically rectangular grids,” *Comp. Math. Applic.*, vol. 33, no. 4, pp. 81–104, 1997.
- [27] J. A. Roden and S. D. Gedney, “Convolution PML (CPML): An efficient FDTD implementation of the CFS-PML for arbitrary media,” *IEEE Microw. Opt. Technol. Lett.*, vol. 27, no. 5, pp. 334–339, 2000.
- [28] J.-P. Berenger, *Perfectly Matched Layer (PML) for Computational Electromagnetics*. Morgan and Claypool, 2007.
- [29] A. Elsherbeni and V. Demir, *The Finite-Difference Time-Domain Method for Electromagnetics with MATLAB Simulations*. Raleigh, NC: Scitech Publishing Inc., 2009.
- [30] J. A. Pereda, L. A. Bielva, A. Vegas, and A. Prieto, “A treatment of magnetized ferrites using the FDTD method,” *IEEE Microw. Guided Wave Lett.*, vol. 3, no. 5, pp. 136–138, 1993.
- [31] J. A. Pereda, L. A. Bielva, M. A. Solano, A. Vegas, and A. Prieto, “FDTD analysis of magnetized ferrites: Application to the calculation of dispersion characteristics of ferrite-loaded waveguides,” *IEEE Trans. Microw. Theory Tech.*, vol. 43, no. 2, pp. 350–357, 1995.
- [32] W. K. Gwarek and A. Moryc, “An alternative approach to FD-TD analysis of magnetized ferrites,” *IEEE Microw. Wireless Compon. Lett.*, vol. 14, no. 7, pp. 331–333, 2004.

- [33] M. Potter and M. Nauta, “FDTD method on an fcc grid for the scalar wave equation and Maxwell’s equations,” *2011 IEEE International Symposium on Antennas and Propagation and USNC/URSI National Radio Science Meeting*, p. 1754, 2011.
- [34] M. E. Potter, M. Lamoureux, and M. D. Nauta, “An FDTD scheme on a face-centered-cubic (FCC) grid for the solution of the wave equation,” *J. Comp. Phys.*, vol. 230, pp. 6169–6183, 2011.
- [35] M. Potter and M. Nauta, “FDTD on a face-centered cubic (FCC) grids for Maxwell’s equations,” *Submitted to IEEE Trans. Antennas Propag.*, 2012.
- [36] M. Nauta, M. Okoniewski, and M. Potter, “FDTD method on a Lebedev grid for anisotropic materials,” *2012 IEEE International Symposium on Antennas and Propagation and USNC/URSI National Radio Science Meeting*, p. 260.7, 2012.
- [37] —, “FDTD method on a Lebedev grid for anisotropic materials,” *Submitted to IEEE Trans. Antennas Propag.*, 2012.
- [38] —, “FDTD on a Lebedev grid for analyzing magnetized ferrites,” in *Electromagnetics in Advanced Applications (ICEAA), 2012 International Conference on*, sept. 2012, pp. 760–763.
- [39] S. Wang and F. L. Teixeira, “Lattice models for large-scale simulations of coherent wave scattering,” *Phys. Rev. E*, vol. 69, p. 016701, 2004.
- [40] B. Finkelstein and R. Kastner, “A comprehensive new methodology for formulating FDTD schemes with controlled order of accuracy and dispersion,” *IEEE Trans. Antennas Propag.*, vol. 56, no. 11, pp. 3516–3525, 2008.
- [41] C. Wagner and J. Schneider, “On the analysis of resonators using finite-difference time-domain techniques,” *IEEE Trans. Antennas Propag.*, vol. 51, no. 10, pp.

2885–2890, 2003.

Appendix A

Waveforms

Excitation waveforms are often implemented naively in FDTD simulations. Dipole sources are a quick and easy excitation that are often sufficient to get the job done. However, such excitations will introduce artifacts that appear as noise in the simulation. These artifacts can be minimized by choosing an excitation that fits the problem type. For example, the excitation should not have frequency content above the Nyquist limit, nor should the spatial distribution excite modes that cannot be accurately modeled with the chosen gridcell parameters.

A.1 Time Domain Waveforms

There are three main concepts that should be understood about the time domain waveform of an excitation. First, the bandwidth of the excitation should be matched to the measurements taken from the simulation. Second, the source should not deposit charge unless the charge buildup has significance to the simulation. Third, excitation waveforms typically decay exponentially to zero as $t \rightarrow \pm\infty$. Therefore, the excitation cannot be exactly reproduced and must be turned on. The magnitude of the unit step at turn on needs to be minimized because it excites frequencies above the desired bandwidth. This section describes a couple of commonly used waveforms and how to ensure that they have an appropriate bandwidth for the simulation.

The excitation used in this thesis is the first derivative of a Gaussian in the time domain:

$$j(t) = A\sqrt{2e}\left(\frac{t-t_0}{\tau}\right)\exp\left(-\left(\frac{t-t_0}{\tau}\right)^2\right) \quad . \quad (\text{A.1})$$

The frequency content is given by

$$|J(f)| = |A|\sqrt{2e\pi\tau}|\pi\tau f|\exp\left(-(\pi\tau f)^2\right) \quad . \quad (\text{A.2})$$

τ can be found from the frequency domain function by setting a max frequency at the 95% bandwidth of the pulse. Once τ has been determined, t_0 can be found such that minimum unit step happens at time zero. In practice it has been found that the unit step should have an amplitude of $10^{-8}j_{\max}$. Figure A.1 shows how τ and t_0 are found by plotting the normalized waveform in frequency and time:

$$\tau = \frac{2.146}{\pi f_{95}} \quad , \quad t_0 = 4.559\tau \quad . \quad (\text{A.3})$$

Another commonly used waveform is known as the Ricker wavelet. It is simply the second derivative of a Gaussian pulse:

$$j(t) = A\left(1 - 2\left(\frac{t-t_0}{\tau}\right)^2\right)\exp\left(-\left(\frac{t-t_0}{\tau}\right)^2\right) \quad , \quad (\text{A.4})$$

$$|J(f)| = 2|A|\sqrt{\pi\tau}(\pi\tau f)^2\exp\left(-(\pi\tau f)^2\right) \quad . \quad (\text{A.5})$$

Figure A.2 shows that

$$\tau = \frac{2.397}{\pi f_{95}} \quad , \quad t_0 = 4.711\tau \quad . \quad (\text{A.6})$$




Figure A.1: Frequency content (top) and time domain waveform (bottom) of an excitation that is the first derivative of a Gaussian. The plots show how to determine the bandwidth and minimum time offset of the pulse. Note that the pulse has no DC content and will not deposit charge at the excitation location.

Both the first and second derivative Gaussians have frequency content that increases from DC, peaks, and then exponentially decreases towards zero. In certain situations, it is advantageous to excite a narrower band of frequencies. This can be done with a cosine modulated Gaussian:

$$j(t) = A \cos(2\pi f_c(t - t_0)) \exp\left(-\left(\frac{t - t_0}{\tau}\right)^2\right) \quad . \quad (\text{A.7})$$

The frequency content of a Gaussian pulse is centered at zero but the cosine shifts it out to a desired center frequency (f_c):

$$|J(f)| = |A|\sqrt{\pi\tau} \exp\left(-(\pi\tau(f - f_c))^2\right) \quad . \quad (\text{A.8})$$




Figure A.2: Frequency content (top) and time domain waveform (bottom) of a Ricker wavelet. The plots show how to determine the bandwidth and minimum time offset of the pulse. Note that the pulse has no DC content and will not deposit charge at the excitation location.

τ and t_0 are determined in the same way as before from the Gaussian envelope. Assuming that the frequency shift is significantly greater than the center frequency ($f_c > f_{95}$) the DC content of the signal is negligible. Figure A.3 shows that

$$\tau = \frac{1.731}{\pi f_{95}} \quad , \quad t_0 = 4.108\tau \quad . \quad (\text{A.9})$$

There is one important caveat with the cosine modulated Gaussian. One of the biggest advantages of time domain simulations is that wideband measurements can be made with a single simulation run. For example, by using the post-processing techniques in Appendix B. Furthermore, the narrower band excitations require longer time offsets and thus, longer simulation lengths. That said, a particular application of the cosine modulated Gaussian is to excite transmission lines above their cutoff frequencies.




Figure A.3: Frequency content (top) and time domain waveform (bottom) of a cosine modulated Gaussian used for narrow band excitations. The plots show how to determine the bandwidth and minimum time offset of the pulse. The DC content of the waveform goes to zero when the center frequency is higher than the pulse bandwidth.

A.2 Spatial Extent of Excitations

The frequency content of a waveform does not change much from simulation to simulation. On the other hand, the spatial distribution of a waveform depends a lot on the task at hand. There are times when an infinitesimal dipole is right for the simulation, for example, when exciting passive antenna elements. However, there are also a lot of cases where such a localized source will make post-processing simulation data unnecessarily complicated. For example, a localized source in a transmission line excites high order modes that are not of practical interest nor accurately sampled by the grid. Alternatively, it is common to assume that a single transmission line mode is used to excite circuit elements. These elements might generate some unwanted modes. In order to measure mode to mode coupling it is necessary to excite a single




Figure A.4: Spatial distribution (top) and excitation amplitude of different modes (bottom) for a sine windowed Gaussian (solid) and an infinitesimal dipole (dotted). The parameters of the sine windowed Gaussian are $x_0/L_x = 0.2$ and $\sigma/L_x = 0.3$.

mode.

Consider the task of exciting the first few modes in a rectangular waveguide where the amplitude distribution of each mode is not yet known. The metal boundaries on each edge will force the tangential components of the electric field to zero. Therefore, the source will be setup such that all field components go to zero at the edge of the boundary. Second, it is reasonable to assume that the solution will be smooth and that modes may have even or odd symmetry. Therefore, the spatial distribution should be neither even nor odd.

A spatial distribution that fits these criteria is an off-center Gaussian windowed by a half sine wave (for a domain from 0 to L_x):

$$j(x) = A(t) \sin\left(\frac{\pi x}{L_x}\right) \exp\left(-\left(\frac{x - x_0}{\sigma}\right)^2\right) . \quad (\text{A.10})$$




Figure A.5: Magnitude of electric field components versus wavenumber at a fixed frequency for a rectangular waveguide. The dipole excitation (top) contains artifacts which are significantly reduced when a spatially distributed excitation (bottom) is used. The peaks are labeled with the FDTD-analytic solution. Note that the TH modes are degenerate with TE modes.

The analytic transform of this function is complicated because the domain only extends from 0 to L_x , but the transform is plotted numerically in Figure A.4. Assuming that only the first 3 modes in each dimension are to be excited (up to TE_{33} for a regular rectangular waveguide). Appropriate values for σ and x_0 are

$$x_0 = 0.2L_x \quad , \quad \sigma = 0.3L_x \quad . \quad (\text{A.11})$$

An xy -plane source distribution should be a function of both x and y .

Figure A.5 shows the importance of using a clean excitation. The rectangular waveguide simulated in Section 4.2.3 was excited using an infinitesimal dipole and a sine windowed Gaussian. The results are plotted with the FDTD-analytic solution[41] to help identify the modes. The waveguide excited by a dipole shows a significant

number of unidentifiable peaks. This is a problem because a key step in extracting the propagation constant versus frequency characteristics is distinguishing the different modes and following the ridgeline of allowed propagation constants as frequency changes. It is much easier to identify modes when the domain is excited cleanly. Note that transients are attenuated in both simulations by starting to record $5L_x$ away from the source.

Appendix B

Extracting Data

B.1 Isolating Propagating Modes

In electromagnetism, the wavelength (λ) of a wave is related to its frequency (f). For plane waves they are inversely proportional to one another but in more complicated situations the relation becomes more complicated. Determining the relationship between wavelength and frequency is extremely important for circuit design. This section describes how to extract the relationship from a single FDTD simulation. The problem is solved by finding the frequency dependence of the propagation constant ($\beta = 2\pi/\lambda$).

It is assumed that the wave propagates in the z -direction and is harmonic in z and t . At a single frequency:

$$\vec{E}(x, y, z, f) = \vec{E}_0(x, y) \exp(j2\pi ft - j\beta(f)z) \quad . \quad (\text{B.1})$$

This is true of propagating modes in lossless material but for evanescent modes or lossy materials the exponential decay must be taken into account. The electric or magnetic field is found over time and 3D space during an FDTD simulation. However, recording every single gridpoint over all time consumes a huge amount of memory. To find the $\beta(f)$ curve it is sufficient to sample along z and t assuming that $\vec{E}_0(x, y) \neq 0$

for the chosen x and y .

The recorded $\vec{E}(z, t)$ is transformed to $\vec{E}(k_z, f)$ using a discrete Fourier transform (DFT). If the domain contains a single allowed mode then the $\beta(f)$ curve is a ridgeline on the surface of $|\vec{E}(k_z, f)|$. However, even if a superposition of modes is present, as is usually the case, the modes can be distinguished from one another by their polarizations. Unfortunately, each mode is excited with a different amplitude and the $\beta(f)$ curves may cross one another. This makes it difficult to fully automate the analysis.

Instead, a graphic user interface (GUI) was developed to extract the curves presented in this thesis. The GUI relies on the following features to produce good results. First, the data is analyzed one frequency at a time. The dot product between an input polarization and the electric field over space is transformed over a narrow range in k -space using the chirp Z-transform. This zoomed in array is superimposed on a plot of the magnitude of the components of the electric field over a larger range of k_z . If the narrow range of data contains the desired peak then k_z at the peak maximum identifies $\beta(f)$ for the frequency being analyzed. The frequency is then incremented. Since $\beta(f)$ is continuous, a simple tracking algorithm automates the next guess. However, the user should verify the guess at each step because the tracker may switch modes if $\beta(f)$ curves cross each other or large neighboring peaks interfere. The process is repeated for each mode found.

B.1.1 Sources of Error

The largest source of error is the Fourier transform leakage error. In continuous problems a resonance peak or $\beta(f)$ at a fixed frequency is represented by a Dirac

delta. However, in a discretized simulation that is propagated for a finite time the Dirac delta must be convolved with a sinc function because the time series is effectively multiplied by a rectangular window. Effectively, all observed peaks have a finite width that is inversely proportional to the simulation time and sample line length.

In some cases there are several modes with similar $\beta(f)$ curves. These modes are difficult to distinguish because of the leakage error. One way of dealing with the problem is to distinguish modes by polarization states. This works especially well for the two orthogonal plane wave modes that have different $\beta(f)$ curves. However, there are other cases where the interfering peaks are not orthogonal.

Other errors arise from the setup of the computational domain. If several modes are being extracted then the spatial distribution of the excitation must excite all of the modes of interest. Ideally with similar amplitudes. In other cases the excitation should be setup to excite a single mode and reduce the effects of leakage error.

B.2 Transmission Line Cross Sections of Propagating Modes

Visualizing the cross section of the fields ($\vec{E}_0(x, y)$) for a propagating mode in a transmission line is also important to circuit designers. For example, excitation feeds and transmission line transitions are designed by matching the fields in one transmission line to the fields in the other. A transmission line cross section also gives information about the source of losses, where to place lumped element components, and whether or not there will be cross talk between neighboring transmission lines.

The challenge with cross section visualization is that the implementation needs to

minimize the amount of data stored during the simulation run. In order to measure $(\vec{E}_0(x, y, k_z, f))$ for any k_z or f , data must be stored over all space and time. However, if a single k_z is chosen then only a plane of points needs to be stored over time. This leads to a significant savings in memory.

Furthermore, a single value of k_z can be computed without using a full DFT along z . Rather, the single value of k_z can be measured using a weighted sum of the field values along each line:

$$\vec{E}_0(x, y, k_z, t) = \sum_{k=0}^{N_z} \vec{E}_0(x, y, k\Delta z, t) (\cos(k_z k\Delta z) + j \sin(k_z k\Delta z)) \quad . \quad (\text{B.2})$$

$\vec{E}_0(x, y, k_z, t)$ is transformed to frequency after the completion of the simulation. The cross section of the desired mode is plotted by determining the frequency at which the mode has the measured propagation constant and plotting $\vec{E}_0(x, y)$ as a quiver plot.

B.3 Measuring Reflection and Transmission

The line sampling method can also be extended to measure the reflection and transmission coefficients of different modes over frequency for a transmission line junction or discontinuity. For example, if plane waves are incident on a slab of dielectric material then the reflection (Γ) and transmission (T) coefficients can be recorded by sampling a line before and after the discontinuity. The Fourier transform automatically distinguishes between forward and backward traveling waves. If two lines of

samples are measured at a fixed x and y position ($\vec{E}_0(x, y, z, t)$, $\vec{E}_1(x, y, z, t)$) then:

$$\Gamma = \frac{\vec{E}_0(x, y, -k_z(f), f) \cdot \hat{p}}{\vec{E}_0(x, y, k_z(f), f) \cdot \hat{p}} \quad , \quad (\text{B.3})$$

$$T = \frac{\vec{E}_1(x, y, k_z(f), f) \cdot \hat{p}}{\vec{E}_0(x, y, k_z(f), f) \cdot \hat{p}} \quad , \quad (\text{B.4})$$

where \hat{p} is a unit vector representing the polarization of the extracted mode. Using the exact polarization state is not necessary for evaluating Γ and T , but it is helpful for removing the interference of neighboring peaks.

Appendix C

Programming Structure

This appendix outlines the object oriented approach to setting up the computational domain and shows how the similarities between the components of the electric and magnetic fields can be exploited. It also shows how polymorphism and inheritance can be used to greatly simplify the recursive loop in code. Adopting these approaches leads to a maintainable software package that is much easier to debug and extend than naive implementations.

First, note that every gridpoint needs to be updated at each timestep regardless of whether that point lies in a bulk material, on a boundary, or in a source region. These different types of updates will be referred to as grid objects. In code, each grid object is its own class because the calculation of the update coefficients depends on its specific purpose. However, all of these grid objects share the same base class because it is assumed all update equations take the same form. The grid object specific calls are only used while setting up the domain. In the recursive update, only the functions of the base class are called. It follows that the spatial limits of the for loop in the recursive update does not depend on the location of any grid objects.

In contrast, many other implementations rely heavily on setting up the spatial for loop for each grid object. These implementations become complicated quickly and are prone to programming errors. Furthermore, it is more difficult to parallelize such




Figure C.1: Structure of the Lebedev grid. The electric and magnetic fields each form an FCC grid but together they form an easily indexed Cartesian grid. All field components are collocated. The update stencil only extends to the immediately adjacent gridpoints of opposite field type.

implementations because the limits of spatial for loops depend on the grid objects the user wants to simulate. If the spatial limits of the loop are determined in advance it is easier to take advantage of hardware architectures and partition memory between different processors.

As justified throughout the thesis, it is assumed that the update equations always take the same form. That is, an update only depends on itself at the previous timestep and the values of the opposite field type at the 6 immediately adjacent gridpoints. This means that each grid object will need to store 7 3×3 matrices of coefficients. With this assumption the size of the stencil is fixed and it can be hard coded into the main recursive update. Therefore, the memory locations that will be accessed are known in advance. As before, this knowledge is extremely beneficial when programming hardware accelerations.

The memory storage scheme can also be setup to make the programming easier

because the Lebedev grid (Figure C.1) uses collocated field components and staggered fields. The gridpoints for each field form an FCC grid which is relatively complex to index. However, together the two grids form a Cartesian grid that is extremely simple to index. Therefore, all of the field components are stored in a single array. This array is indexed by component u and location indices i, j, k . The field type is implicitly given by the grid index. For example, if $i + j + k$ is even then that location represents an electric field. The memory structure means that the fields required for an update will lie at $\text{EH}[i + p\delta_{wx}, j + p\delta_{wy}, k + p\delta_{wz}]$ for $p = \pm 1$ and $w = x, y, z$. Any update that lies on the boundary will attempt an access with a negative index. This issue is dealt with by wrapping negative indices around the array. This solution simultaneously implements periodic boundary conditions if no other boundary condition is specified and the number of cells is a half integer. Note also that it is easy to implement both electric and magnetic field updates in a single for loop ($f = 0, 1$ for E, H) because the fields are located in the same array and the equations take the same form. In contrast, the main loop typically updates all of E, H, D , and B with separate pieces of code.

It is also advantageous to assume that many gridpoints share that same grid object. This is definitely true of all the simulations presented in the thesis. In this case, it is most efficient to store a list of all of the grid object classes used in a simulation and store a single integer reference into the list at every gridpoint in the computational domain. This gives significant memory savings because a single integer is stored at each gridpoint instead of $7 \times 3 \times 3$ matrices.

The computational domain is setup by precomputing the coefficients in the list of grid objects and ensuring that each gridpoint references the appropriate grid object.

The computational domain is initialized with a single background material. Then, grid objects are added to the domain by appending them to the grid object list and changing the reference index at the appropriate grid points. The equations for adding metals and materials to the grid are phrased such that they use information from update coefficients that have already been computed. This is done so that adding objects is easier. If the grid object being added is large then it is possible that it will intersect more than one grid object that is already on the domain. A new grid object is appended for each case. This may seem complicated, but it is well suited to automatic generation of the coefficients. With the naive approach, the programmer is responsible for adding the appropriate update for every case. This is often tedious, for example, at the intersection of two materials in a rectangular waveguide.

Another benefit of the Lebedev grid is that all three components are updated at every point. Therefore, the calculation of the update coefficients can be expressed in any direction by a permutation of the indices. The advantage is that calculations can be written with a single piece of code that is independent of orientation. When debugging and testing, the programmer knows that all implementations work the same. In contrast, copy-and-paste style coding is commonly used to permute the indices. Code written with this technique is difficult to maintain because changes and testing may only cover a subset of the cases.

The grid storage structure, fixed update form, and grid symmetry can be combined to give a very simple main update loop. The pseudo-code for this loop is given below. The fact that this piece of code is small enough to fit on a single page is the whole point of this appendix.

```

for t = 0.0 to maxTime in steps of dt
  for f = 0,1 // for E, then H
    for (i,j,k) = (0..2Nx,0..2Ny,0..2Nz) // for all grid points
      if (i+j+k)%2 == f // even => E, odd => H
        oIndex = ObjectIndex[i,j,k] // reference index
        tmpF = {0,0,0} // field at next time

    // *** Add inv(a)b_uv F[v] terms
    for u = 0,1,2 // update Fx,Fy,Fz
      for v = 0,1,2 // depends on Fx,Fy,Fz
        if b[oIndex](f,u,v) != 0 // inv(a)b_uv != 0
          tmpF[u] += b[oIndex](f,u,v)*EH(v,i,j,k)

    // *** Add inv(a)c[p,w]_uv G[v] terms
    for p = -1,1 // shift by +/-1
      for w = 0,1,2 // shift direction
        for u = 0,1,2 // update Fx,Fy,Fz
          for v = 0,1,2 // depends on Gx,Gy,Gz
            if c[oIndex](f,u,v,p,w) != 0 // inv(a)c[p,w]_uv != 0
              tmpF[u] += c[oIndex](f,u,v,p,w)*
                EH(v,i+p*kDelta(FIELD_X,w),
                  j+p*kDelta(FIELD_Y,w),
                  k+p*kDelta(FIELD_Z,w))

    // *** Add constant terms
    for u = 0,1,2 // update Fx,Fy,Fz
      if A[oIndex](f,u) != 0 // A(t,E,H) != 0
        tmpF[u] += A[oIndex](f,u,t)

    for u = 0,1,2 // update Fx,Fy,Fz
      EH(u,i,j,k) = tmpF[u] // save fields

    t += dt/2.0 // update time

    for oIndex = 0 to numObjects
      update time dependent coefficients as needed (sources)

    update and add PML auxiliary functions

    store fields/output data as required

```

