



KUNGL
TEKNISKA
HÖGSKOLAN

Hybrid Time-Domain Methods and Wire Models for Computational Electromagnetics

Gunnar Ledfelt

Stockholm 2001

Doctoral Dissertation
Royal Institute of Technology
Department of Numerical Analysis and Computer Science

Akademisk avhandling som med tillstånd av Kungl Tekniska Högskolan framläggas till offentlig granskning för avläggande av teknisk doktorsexamen fredagen den 30 mars 2001 kl 13.00 i sal D2, Kungl Tekniska Högskolan, Lindstedtsvägen 5, bv, Stockholm.

ISBN 91-7283-058-1

TRITA-NA-0106

ISSN 0348-2952

ISRN KTH/NA/R--01/06--SE

© Gunnar Ledfelt, March 2001

Universitetsservice US AB, Stockholm 2001

Abstract

The Finite-Difference Time-Domain (FD-TD) method is the most commonly used time-domain method for solving the Maxwell equations. The FD-TD method pioneered by K.S. Yee 1966 is an explicit finite-difference scheme using central differences on a staggered Cartesian grid (Yee grid) and is second-order accurate in both space and time. It has been attractive for industrial users since the early 1980s because the basic method is relatively simple to program and the geometry handling is fairly straightforward. The main drawback of the FD-TD method is its inability to accurately model curved objects and small geometrical features. The Cartesian FD-TD grid leads to a staircase approximation of the geometry and parts smaller than an FD-TD cell might be neglected by the grid generator. We present three different methodologies to minimize this drawback of FD-TD but still benefit from its advantages. They are parallelization, hybridization with unstructured grids, and subcell models for thin wires.

Parallel computers can solve very large FD-TD problems. This is illustrated by a lightning problem for a real aircraft where more than one billion FD-TD cells are used. The cell size is one inch which gives a very fine grained grid. This type of simulation is important for electromagnetic compatibility problems where the complexity of the geometry requires small cells to give accurate results.

The idea behind our hybridization method is to use unstructured boundary fitted grids to resolve the geometrical features and Yee grid elsewhere. In this way we avoid the staircasing problems, and small parts can be resolved but still keep the efficiency of FD-TD. In two dimensions, our hybrid methods are second-order accurate and stable. This is demonstrated by extensive numerical experiments. In three dimensions, our hybrid methods have been successfully used on realistic geometries such as a generic aircraft model. The methods show super-linear convergence for a vacuum test case and almost second-order convergence for a perfect electric sphere. However, they are not second-order accurate. This is shown to be caused by the interpolation needed when sending values from the Yee grid to the unstructured grid. Stability issues are also discussed.

The cross-section of thin wires are smaller than the Yee cells and hence subcell models for thin wires have been developed for FD-TD. We present a new model for arbitrarily oriented thin wires. Previously published models for FD-TD require the wire to be oriented along the edges of the grid and hence a staircasing error is introduced. The new model avoids these errors. Results are presented illustrating the superiority of the new thin-wire subcell model.

ISBN 91-7283-058-1 • TRITA-NA-0106 • ISSN 0348-2952 • ISRN KTH/NA/R--01/06--SE

Contents

1	Introduction	3
1.1	Computational Electromagnetics	3
1.2	Overview and main results	4
1.3	The PSCI project GEMS	9
1.4	List of papers	10
1.5	Division of work	11
2	A Popular Description in Swedish	13
2.1	Att lösa Maxwells ekvationer numeriskt	13
3	The Maxwell Equations	21
3.1	The equations	21
3.2	Reduction to two dimensions	23
3.3	Reduction to one dimension	23
3.4	Integral formulation	23
3.5	The wave equation	24
3.6	Material properties	25
4	FD-TD	27
4.1	Introduction to FD-TD	27
4.2	Discretization used in FD-TD	27
4.3	The leap-frog scheme	30
4.4	Stability conditions	31
4.5	Performance of the leap-frog update	31
4.6	Boundary conditions	33
4.7	Sources	36
4.8	Visualization as debugging and validation tool	36
4.9	Parallelization	37
4.10	Subcell models	38
4.11	Near-to-far-field transformations	38
4.12	Frequency-dispersive materials	39
4.13	Divergence-free nature	39

5	Large Scale FD-TD — A Billion Cells	41
5.1	Introduction	41
5.2	Parallel implementation	42
5.3	One billion cells	43
5.4	Visualization of large FD-TD data	45
5.5	Conclusions	47
5.6	Acknowledgments	47
6	Hybrid Methods in 2D	49
6.1	Introduction	49
6.2	Finite-difference method	49
6.3	Finite-element method	50
6.3.1	FE-TD formulation	50
6.3.2	Spatial discretization	51
6.3.3	Time discretization	53
6.3.4	Workload and memory requirements	56
6.4	Finite-volume method	56
6.4.1	FV formulation	56
6.4.2	Spatial discretization	56
6.4.3	Time discretization	60
6.4.4	Preservation of divergence	60
6.4.5	Creating the dual grid	62
6.4.6	Workload and memory requirements	63
6.4.7	Stability analysis	65
6.5	Grid requirements	67
6.6	Hybridization	69
6.6.1	FD-FE	71
6.6.2	FD-FV	71
6.6.3	FE-FV	72
6.7	Stability	73
6.7.1	Details	74
6.7.2	FD-FE	74
6.7.3	FD-FV	75
6.8	Convergence	75
6.8.1	Modeling of circular cylinders in FD-TD	76
6.8.2	The coarse grids	76
6.8.3	Grid refinement	78
6.8.4	Results	80
6.9	PMC wall	86
6.9.1	Details of the numerical setup	86
6.9.2	Results	87
6.10	Conclusion	89

7	Hybrid Methods in 3D	91
7.1	Introduction	91
7.2	Finite-difference method	91
7.3	Finite-element method	92
7.3.1	Huygens' surfaces for the FE solver	93
7.4	Finite-volume method	96
7.4.1	Space discretization	96
7.4.2	Time discretization	99
7.4.3	Preservation of divergence	99
7.4.4	Stability analysis of the FV solver	100
7.5	Hybridization	102
7.6	Stability of the hybrid methods	104
7.7	Results	107
7.7.1	Convergence in vacuum	107
7.7.2	Scattering from a PEC sphere	110
7.7.3	Scattering from a dielectric sphere	113
7.7.4	Scattering from the NASA Almond	114
7.7.5	Scattering from the generic aircraft <i>RUND</i>	115
7.8	Conclusions	118
8	Thin-Wire Subcell Models	119
8.1	Introduction to thin-wire subcell models	119
8.2	Physics	121
8.2.1	First assumption	123
8.2.2	Second assumption	124
8.3	The wire equations	127
8.4	Coupling between the wire and the surrounding 3D domain	127
8.4.1	Distributional contemplation	128
8.4.2	Distributed coupling between the wire and the surrounding 3D domain	129
8.4.3	Effects of distribution	130
8.5	Discretization of the wire equations	131
8.5.1	Inner nodes	132
8.5.2	Boundary conditions	133
8.6	Discrete distribution	134
8.6.1	Introduction to discrete interpolation	134
8.6.2	Discrete interpolation for the wire problem	137
8.7	Improved discrete distribution	138
8.8	Stability	140
8.9	Comments on the physical assumptions	145
8.10	Results	146
8.10.1	Receiving dipole antenna	147
8.10.2	Transmitting dipole antenna	155
8.10.3	Loop antenna	161
8.10.4	The wire length	163

8.11 Conclusion	165
9 Color Electromagnetics	167

Acknowledgments

A thesis is a one-man-work but it would be deeply unfair not to recognize some people whom have had a profound impact on my thesis.

First I would like to thank Ulf Andersson for being my best friend and colleague on the journey of CEM exploration. Without you the course would have been more winding and much lonelier. I really appreciate all the moments when we solved problems on the white-board and I am grateful for all your patience shown at these discussions.

Sincerest appreciation is due to my supervisor Prof. Björn Engquist for giving me most valuable directions on the wave-propagation subject. Your strong belief in academic gain in industrial collaboration gave me the opportunity to work in the GEMS project. At several occasions you also arranged very valuable visits for me at UCLA which I am truly thankful for.

Fredrik Edelvik, Lasse Eriksson, Lennart Hellström and Erik Abenius, special thanks to all of you for the fun we had during the development of the codes. Fredrik, for all the joy in the development of the Finite-Volume solver. Lasse, for the thrilling moments of running the first hybrid simulations with the Finite-Element solver. Lennart, for all your help with the processing of the hideous input-files from the end-users. Erik, for the stimulating start on the thin wire work.

I thank Bo Strand for leading the PSCI-part of the GEMS project and Anders Ålund for teaching me good programming style. Thanks to all industrial partners who put my work into a much broader perspective. Among all of you, Jonas Gustafsson always pushed me harder. You forced me, without mercy, into the obstacle of reality where spheres and cylinders are rare. You also let me understand that there is an industrial need for my results which was indeed inspiring. Thank you.

I am very grateful to Prof. Jesper Ooppelstrup. You have spent much time explaining trivial and difficult matters. Also, the inspiration you have given me is truly appreciated.

Daniel Noreland and Kurt Otto, thank you for all the fruitful discussions we had during both office hours and after sunset.

I would also like to thank Henrik Holter. Your pedagogic and patient explanations regarding the physics of electromagnetics have been very enlightening.

My UNIX skills are mostly due to Harald Barth, Kjell Högström and Janne Kärroman. Thanks for all your magical UNIX-scripts making life much easier.

I am grateful to Per Ekman and Erik Engquist who joined me into 3D immersive VR visualization. I solved many problems when I gazed into the nodes and beams of 3D dual grids. Thanks for all VR help.

A big thanks also to my colleagues at NADA and TDB. We have had a joyful time together.

My scientific career started many years ago when my father taught me to pose the question “Why is it like that?”. This is still one of my most used questions. Thank you for enticing me into mathematics and physics.

Last but not least I would like to thank Ninni for all your love, support and patience. You have always showed interest in my research and very seldom asked when (never if) I will finish this thesis.

Financial support has been provided by NADA, KTH, NUTEK (PSCI), FRN (PDC), NFFP, Knut och Alice Wallenbergs stiftelse, Telefonaktiebolaget L M Ericssons stiftelse för främjandet av elektroteknisk forskning, Frans Georg och Gull Liljeroths Stiftelse, and Ragnar och Astrid Signeuls fond,

Going through this acknowledgment might lead you to the erroneous conclusion that this work is not my own but the work of my colleagues. If that is the case, please ask me to prove my skills by writing another thesis without the help of the people above! But not on the topic of a 3D hybrid Maxwell solver for industrial use. ;-)

To Ninni

Chapter 1

Introduction

1.1 Computational Electromagnetics

The Maxwell equations were formulated in 1873 by James Clerk Maxwell [66] and describe electromagnetic wave phenomena. Electromagnetic waves are frequently used for many purposes, for example radio waves, micro waves and radar waves. During the last few decades applications for electronic devices have found their way into many different disciplines such as medicine, telecommunication, transportation systems and consumer electronics. In many of these cases a thorough understanding of the electromagnetic behavior is fundamental and hence there has been a lot of effort put into solving the Maxwell equations. There is only a few rare cases when the Maxwell equations can be solved exactly. Therefore approximate mathematical methods and measurement techniques have been developed. When computers became available a third class of methods started to evolve; numerical methods.

Development and analysis of numerical methods for solving the Maxwell equations are today a very active field of research. This might seem strange because the equations are in fact quite simple. Furthermore, boundary conditions and interface conditions are also simple to formulate. But the combination of the equations and the boundary/interface conditions, together with the oscillatory nature of the electromagnetic fields, constitute big challenges.

Numerical methods for the Maxwell equations, or computational electromagnetics (CEM), becomes more and more recognized in industry and some important applications for CEM include:

- Antenna problems
- Electromagnetic compatibility
- Microwave systems
- Dosimetry
- Radar cross section predictions

1.2 Overview and main results

This section gives an overview of the chapters in this thesis. Chapters 1–4 contain known material but forms a necessary foundation for Chapters 5–9 where new material is presented.

Chapter 1, Introduction

A short introduction to the topic of the thesis is given in this Chapter. A summary and main results are also given here. The research and development described in this thesis have been conducted in the GEMS project within the Parallel and Scientific Computing Institute (PSCI). GEMS stand for General Electromagnetic Solvers and the project is briefly put into its context in Section 1.3.

The thesis is partly based on a number of papers which are listed on Page 10. My contribution to the collaborative parts of the work presented is described in Section 1.5.

Chapter 2, A Popular Introduction in Swedish

In Chapter 2 a popular introduction to electromagnetic wave problems is given. It is focused on applications rather than methods but a short presentation of the Finite-Difference Time-Domain method (FD-TD) is given. The drawbacks of the FD-TD method are illuminated and the need for hybrid methods is indicated. This Chapter is written in Swedish and intended for the general public.

Chapter 3, The Maxwell equations

The Maxwell equations given in (3.1) are general but for certain classes of problems the equations can, or must, be simplified.

Let us define the electric size as the relation between geometric length scales and wavelengths (see Figure 1.1). For a problem where the interesting wavelengths are several orders of magnitude longer than the geometric length scale, the object is said to be electrically small. On the other hand, an object with a geometric size that is several orders of magnitude larger than the interesting wavelengths is regarded as electrically large.




Figure 1.1. The relation between wavelength and object size defines the electric size of an object.

For electrically small problems the wave nature of the Maxwell equations is generally dominated by electric or magnetic conductive currents. In these cases the time derivative of the fields are small and the governing equations becomes elliptic in character. The finite-element method have in these cases been prosperously used for many years

For electrically large objects, we can solve the Maxwell equations asymptotically for certain canonical geometries. In 1962 Keller [53] showed how these asymptotic solutions, for certain canonical geometries, can be combined to find solutions for more complex geometries. The method is called geometrical theory of diffraction (GTD) and there are several refined methods developed, for example the uniform theory of diffraction (UTD) [55]. This kind of approximations are necessary because solving the full Maxwell equations requires a discretization of several points per wavelength. Hence, an electrically large problem cannot be sufficiently discretized using the computers of today.

We will restrict ourselves to problems of medium electric size. These problems are often solved in either the frequency domain (FD) or the time domain (TD). In FD we solve the problem for a fix frequency and in TD we follow the evolution of wave pulses using a timestepping mechanism. Numerical methods for medium electric sized problems can be applied either directly on the Maxwell equations to yield a volumetric problem or applied to boundary integral formulations to give a problem restricted to surfaces. In this thesis we present methods for solving the Maxwell equations in time domain using volumetric methods.

Chapter 4, FD-TD

The Finite-Difference Time-Domain (FD-TD) method is the most commonly used time-domain method for solving the Maxwell equations. It was introduced by Yee in 1966 [110] and is often referred to as the Yee scheme. The FD-TD method is an explicit finite-difference scheme using central differences on a staggered Cartesian grid (both space and time), i.e. it is a leap-frog scheme. It is second-order accurate in both space and time. One cell of a 3D FD-TD grid is given in Figure 1.2 where the staggering of the six components is shown. The electromagnetic variables located on the other edges/surfaces belong to neighboring cells and are not included in the figure.

Several books have been published dealing with the FD-TD scheme [56, 98, 99, 52, 97]. The survey paper by Shlager and Schneider that appeared in [99] illustrates the rapid growth in the use of FD-TD. Two papers were published during the 1960s and the growth rate has been roughly tenfold per decade since the first publications.

The FD-TD method has been attractive for industrial users since the early 1980s because the basic method is relatively simple to program and the geometry handling is fairly straightforward. The method can also be efficiently implemented on vector computers which made it feasible to solve complex problems on the early supercomputers. As an example, in 1987 SAAB performed lightning analysis on the Swedish fighter aircraft Gripen on a grid with approximately $60 \times 30 \times 30$ cells.




Figure 1.2. Positions of the electric and magnetic field vector components in a unit Yee cell.

Chapter 5, Large Scale FD-TD — A Billion Cells

In this Chapter, Paper 6 listed on Page 10 is reprinted. The paper demonstrates the possibility to solve enormously large FD-TD problems by using parallel computing. It also shows that it is possible to achieve perfect scale-up on a parallel computer with reasonably fast communication, e.g. an IBM SP, for industrially relevant FD-TD simulations

In the paper a lightning simulation on the SAAB 2000 aircraft is presented where more than one billion cells were used. This was run on 125 nodes of an IBM SP. The performance of the code was 20–25 Gflop/s which is very good considered that the peak performance of the nodes was 80 Gflop/s

The billion cell simulation was performed during the summer of 1998. It was presented at SuperComputing-98, October 1998 in Orlando USA, as an invited demonstration at the IBM booth.

To our knowledge this was the first FD-TD simulation published using more than a billion cells and the paper was also acknowledged by Miller in [67].

Today there are many industries, national agencies and academic sites that can make a billion cell FD-TD simulation. If the largest computer of today, "ASCI White", could be allocated for a similar FD-TD computation a hundred times larger simulation could be performed.

Chapters 6 and 7, Hybrid methods in 2D and 3D

The main drawback of the FD-TD method is its inability to accurately model curved objects and small geometrical features. The Cartesian FD-TD grid leads to a staircase approximation of the geometry and parts smaller than an FD-TD cell might be neglected by the grid generator.

There are several suggested remedies to circumvent the effects of the errors introduced by the staircase approximations. They include Cartesian subgridding [79],

conformal modeling [21], multiblock body fitted structured grids [92], unstructured grids [59] and overlapping grids [111].

Another approach to avoid staircasing is to use unstructured grids near curved objects and around small geometrical details but revert to structured grids as quickly as possible for the rest of the computational domain. This combines the efficiency of structured grids with the flexibility of unstructured grids. Wu and Itoh [108] were first to present a combination of the FD-TD method and an implicit finite-element method (FE). They have been followed by several others [71, 54, 90, 112]. A combination of explicit finite volumes (FV) and FD-TD have been proposed by Riley and Turner [84].

In Chapters 6 and 7 we present new techniques to hybridize structured FD-TD solvers and unstructured solvers. The optimal choice of method on the unstructured grid depends on the cell sizes. An implicit method is preferable for grids with small cells while an explicit method is preferable for grids where the cell sizes are of the same order as in the Cartesian grid. Hence, our approach is to implement both an explicit FV solver and an implicit FE solver and couple these solvers with the FD-TD solver.

The hybridization is performed over a transition layer of cells that coexists in both the structured grid and the unstructured grid. Rectangles or rectangular parallelepipeds (2D or 3D) constitute the structured cells in the transition layer whereas triangles and tetrahedra form the unstructured cells in the transition layer.

In Chapter 6 we demonstrate that staircasing of a circular cylinder destroys the second-order accuracy of FD-TD and that second-order accuracy is recovered when using our hybrid methods. This is shown for a perfect electric conducting cylinder, a perfect magnetic conducting cylinder, a dielectric cylinder and diamagnetic cylinder. A vacuum problem is also investigated in order to examine reflections from the transition layer and these errors are also of second-order accuracy. The accuracy is further demonstrated for a 45 degrees inclined perfect magnetic conducting wall where our hybrid technique outperforms two contour path modeling schemes.

Numerical stability experiments are presented where it is found that the FD-FV hybrid is stable as long as a CFL condition is not violated. We do encounter stability problems for the FD-FE hybrid for some grids when the Crank-Nicholson scheme is used for time discretization. However, using the backward differentiation formula (BDF-2) as timestepping scheme regains stability. Making the Crank-Nicholson scheme slightly more implicit does also regain stability without affecting accuracy significantly.

In Chapter 7 the 3D hybrid technique is presented. Results are supplied showing that industrial relevant scattering problems can be successfully solved.

The order of accuracy is studied for a PEC sphere and a vacuum problem. For the PEC sphere almost second-order accuracy is obtained for the discretizations used but for the vacuum problem the order of accuracy is only superlinear. This is caused by an interpolation procedure necessary in the transition layer. However, the accuracy is demonstrated for the NASA almond model problem where the errors from the transition layer are small compared to the signal levels of the PEC

reflections. Good results are also obtained for the generic aircraft *RUND* indicating that errors emanating from the transition layer are not ruining the solution.

The implicit FE-TD solver is unconditionally stable as a stand alone solver [60]. The explicit FV-TD solver, as a stand alone solver, is stable on orthogonal grids as long as the time step is chosen properly. On general grids a node filter is applied to suppress the amplitude of the highest frequency components. In practice, for scattering problems, this filter improves the stability without losing accuracy in the solution. However, in terms of stability, cavity problems are in general much more demanding than scattering problems due to longer simulations that could require millions of timesteps. Methods to improve the stability of the FV-TD solver for resonant cavities are currently under investigation.

We demonstrate in Chapter 7 that our hybrid solvers produces good results for scattering problems but for highly resonant cavities our hybrid solvers eventually become unstable. This happens even if the unstructured grid is orthogonal in which case all stand alone solvers are stable. Hence, these instabilities are caused by the hybridization.

The use of pyramidal cells in the transition layer described in [88] is currently investigated in order to restore the stability for the FD-FE hybrid.

In Chapter 7 we also describe how to incorporate plane wave excitation in the FE solver. This functionality makes it possible to include port excitation, a very important feature for the industrial partners within the GEMS project.

Chapter 8, Thin-Wire Subcell Models

In this Chapter we present a new and accurate method to deal with the problem of how to include thin wires into the FD-TD method. Thin wires are often important parts in electromagnetic compatibility problems and antenna problems. The most straightforward technique to include wires into the FD-TD method would of course be to model the wire as a perfect electric conductor, and hence set all tangential electric components on the wire surface equal to zero. But this strategy requires a very fine discretization in order to resolve the cross-section of the wire which is clearly out of question in most practical cases.

In this Chapter a thin-wire model for FD-TD is extended to allow for an arbitrary orientation of the wire. This is important in order to treat wires that cannot be aligned with the Cartesian grid, for example circular loop antennas and tilted straight wires.

There are other techniques to solve problems including thin wires than using subcell models in FD-TD. These are discussed in Section 8.1. But for volumetric time-domain solvers there are basically two subcell models developed, the model of Umashankar et al. [103] and the model of Holland et al. [49]. These models are both limited to treat wires oriented along the Cartesian grid directions and to our knowledge there are no published papers describing how to include arbitrarily oriented thin wires into an FD-TD simulation.

In this Chapter we present an extension to the subcell model of Holland et al. [49] that gives very good results for circular loop antennas and tilted straight wires. The extension utilizes an innovative interpolation technique based on tri-linear interpolations of the surrounding tangential electric field to several points on a circular shell around the wire. The field on the shell is then averaged to yield the electric field driving the wire equations. The coupling from the wire equations back to the surrounding 3D field utilizes a reversed process.

Results are presented which clearly illustrate the success of the new model. Without the new model, staircasing is the only alternative to discretize arbitrarily oriented thin wires in FD-TD. In some cases the staircasing gives errors of several tens of percent whereas the results for the circular loop antennas and tilted straight wires are confined within one percent deviation.

Chapter 9, Color Electromagnetics

This Chapter discuss scientific visualization of results from computational electromagnetics. It points out some fundamental elements of visualization to enhance the perception of 3D data and 3D geometries. Six pages of color images are included to underline the importance of taking these elements into account when producing color plots.

1.3 The PSCI project GEMS

The research and development described in this thesis have been conducted in the GEMS project within PSCI. The Parallel and Scientific Computing Institute, PSCI [81], was created 1995 to improve the interaction between academia and industry and strengthen Swedish efforts in industrial applications of high performance computing. PSCI is a center of excellence funded by an industrial consortium, NUTEK [78], KTH and Uppsala University. The industrial partners involved include large Swedish companies as well as smaller enterprises.

The GEMS project is one of the larger projects within PSCI. GEMS stands for General ElectroMagnetic Solvers and is a code development project started 1998 and ending spring 2001. The main objective of the GEMS project is to develop a software suite for the Maxwell equations. This software suite aims to be state of the art of the international level and form a platform for future development by Swedish industry and academia.

The core of the software suite is two hybrid codes, one for the time domain (TD) and one for the frequency domain (FD). The TD code is a multi-block and out of core solver based on Finite Differences on structured grids, explicit Finite Volumes and implicit Finite Elements on unstructured grids. The FD hybrid code is based on the Method of Moments, Physical Optics and Geometrical Theory of Diffraction. The coupling between the different FD solvers is taken into account by using iterative techniques and the linear system of equations are solved by iterative

solvers. A Fast Multipole method is also developed to boost the performance of the FD code.

The industrial partners in GEMS are Ericsson Microwave Systems, Saab Ericsson Space and Ericsson Saab Avionics. Code developers are PSCI, the Swedish Institute of Applied Mathematics and the Swedish Defense Research Establishment. Project funding is provided not only by the participants but also to large extent by the National Aeronautical Research Program (NFFP) and the total project work load is more than 400 man months where PSCI conducts approximately 300 man months. Typical projects in the PSCI research programs are joint undertakings where application experts from the industrial partners, graduate students and senior scientists contribute. The criteria for a project include motivation by industrial needs, and theoretical development suitable for PhD thesis work. A direct result of the GEMS project is this thesis which to a large extent is influenced by the joint work between industry and academia.

1.4 List of papers

This thesis is partly based on material from the following papers:

1. Gunnar Ledfelt, Fredrik Edelvik, Lasse Eriksson, and Ulf Andersson. Hybrid time domain solver for the Maxwell equations. In Jan Olov Gustafsson, editor, *RadioVetenskap och Kommunikation 99*, pages 70–74, 371–379 Karlskrona, Sweden, June 1999. SNRV, NUTEK, Högskolan i Karlskrona/Ronneby.
2. E. Abenius, U. Andersson, F. Edelvik, L. Eriksson, and G. Ledfelt. Hybrid time domain solvers for the Maxwell equations in 2D. Technical Report 00:01, PSCI, Parallel and Scientific Computing Institute, KTH, SE-100 44 Stockholm, Sweden, February 2000. Available at <http://www.psci.kth.se/Activities/Report>
3. Fredrik Edelvik, Gunnar Ledfelt. Explicit Hybrid Time Domain Solver for the Maxwell Equations in 3D. *J. Sci. Comput.*, vol 15, nr 1, 2000.
4. Gunnar Ledfelt, Fredrik Edelvik and Ulf Andersson. Hybrid Time Domain Solver for the 3D Maxwell Equations. In U. Zander, editor, *ANTENN 00 – Nordic Antenna Symposium*, pages 57–62, Lund, Sweden, September 2000. FMV, SNRV.
5. Gunnar Ledfelt. A thin wire sub cell model for arbitrary oriented wires for the FD-TD method. In G. Kristensson, editor, *EMB 98 – Electromagnetic Computations for analysis and design of complex systems*, pages 148–155. SNRV, November 1998.
6. Ulf Andersson and Gunnar Ledfelt. Large scale FD-TD—A billion cells. In *15th Annual Review of Progress in Applied Computational Electromagnetics*, volume 1, pages 572–577, Monterey, CA, March 1999.

1.5 Division of work

- The material in Chapters 1 and 2 is completely my own.
- Chapters 3 and 4 were written with Ulf Andersson. They also appear in his Ph.D. thesis [5].
- Chapter 5 is identical to Paper 6 in the list in Section 1.4 where Ulf Andersson was responsible for the parallelization of the FD-TD solver as well as the performance evaluation, and I was responsible for the geometry manipulation and the visualization.
- The work presented in Chapter 6 is a collaborative effort. It is based on Papers 1 and 2 in the list in Section 1.4. The contributions of each author were: Erik Abenius was responsible for the 2D FD-TD implementation, Ulf Andersson was responsible for the design and realization of the numerical evaluation, Fredrik Edelvik designed and implemented the FV-TD solver, Lasse Eriksson designed and implemented the FE-TD solver, and I designed and implemented the hybridization strategy. I was also co-developer regarding the unstructured solvers and the numerical evaluation. Chapter 6 is also part of Ulf Andersson's Ph.D. thesis [5].
- Chapter 7 also describes a joint effort, partly based on Papers 3 and 4 in the list in Section 1.4. The contribution of the authors were: Ulf Andersson was responsible for the 3D FD-TD implementation, Fredrik Edelvik designed and implemented the FV-TD solver, Lasse Eriksson designed and implemented the FE-TD solver, and I designed and implemented the hybridization strategy. I was also co-developer regarding the three separate solvers. The design and realization of the numerical evaluations were shared between Ulf Andersson, Fredrik Edelvik and myself.
- The material in Chapters 8 and 9 is completely my own, with the exception of Figures 9.7 and 9.8 which were produced by Jonas Gustafsson at Ericsson Saab Avionics.

Chapter 2

A Popular Description in Swedish

2.1 Att lösa Maxwells ekvationer numeriskt

‘I begynnelsen skapade Gud himmel och jord. Jorden var öde och tom, djupet täcktes av mörker och en gudsvind svepte över vattnet. Gud sade “ljus, bli till!” Och ljuset blev till. Gud såg att ljuset var gott, och han skiljde ljuset från mörkret. Gud kallade ljuset dag, och mörkret kallade han natt. Det blev kväll och det blev morgon. Det var den första dagen.’ (Första Moseboken 1:1-5)

Orden, som är hämtade ur första kapitlet i Bibeln, ger oss en uppfattning om ljusets betydelse i tillvaron. Ljuset är en typ av elektromagnetisk strålning, precis som radiovågor och värmestrålning. Flera viktiga egenskaper hos ljuset har varit kända under mycket lång tid. Även magnetism och elektricitet har varit känt i tusentals år. Redan under antiken kände man till begreppen. Ordet magnet härrör från ett område i Thessalien i Grekland som heter Magnesia. Där hittade man ett mineral, som nu kallas Magnetit, som attraherar järnföremål. Det grekiska ordet för bärnsten, elektron, har gett namn åt de fenomen som uppträder om man gnider bärnsten mot tyg. Lätta strån och papper dras mot bärnstenen och små sprakande urladdningar i form av gnistor kan observeras. Men det var först under 1700- och 1800-talet som man systematiskt började undersöka och beskriva elektriciteten och magnetismen.

Det skulle dröja ända till 1864 innan James Clerk Maxwell kunde beskriva elektromagnetisk strålning ordentligt. På basis av Faradays beskrivning av det elektriska fältet kunde Maxwell formulera den uppsättning ekvationer som idag bär hans namn. 1873 publicerades Maxwells teori som beskriver ljuset som en elektromagnetisk vågrörelse, där elektriska och magnetiska fältstyrkan periodiskt varierar i plan som är vinkelräta mot fortplantningsriktningen. Teorin angav dessutom att

fortplantningshastigheten skulle vara 300 000 km/s. Den blev dock allmänt accepterad först när Hertz 1888 kunde visa att strålningen från en resonant elektrisk krets hade samma utbredningshastighet som ljuset och dessutom uppvisade andra karaktäristiska ljusegenskaper.

Betydelsen av att kunna förstå och lösa Maxwells ekvationer är idag viktigare än någonsin. I vårt samhälle är vi omgivna av maskiner och apparater som på ett eller annat sätt utnyttjar de fenomen som beskrivs av teorin som Maxwell formulerade för drygt hundra år sedan. Vi ser på TV, lyssnar på radio och pratar i mobiltelefon dagligen. Vi åker bil, tåg och flyg utan att ägna Maxwell en tanke.

Det finns många olika tillämpningsområden för elektromagnetisk analys. Ett problemområde är EMC, electromagnetic compatibility. EMC handlar om att en apparat skall fungera i den elektriska miljö den befinner sig i och den skall heller inte sända ut oönskade elektromagnetiska vågor eftersom dessa kan störa andra apparater. EMC-området blir mer allt viktigare ju snabbare, mindre och strömsnålare de elektroniska apparaterna blir. Exempel på EMC-tillämpningar är sjukhuselektronik och fordons elektronik som inte får störas av blixnar, mobiltelefoner och annan elektronik.



Figur 2.1. På många sjukhus är det förbjudet att använda mobiltelefoner eftersom elektromagnetiska vågor kan äventyra livsviktig sjukhusutrustning.



Figur 2.2. Personlig elektronisk utrustning får inte användas under start och landning när man flyger eftersom det finns risk att flygplanselektroniken störs.



Figur 2.3. Elektroniken i moderna bilar måste skyddas så att inte krockkudden löser ut när mobiltelefonen används.



Figur 2.4. Om en båt med stark radar passerar ett tåg på en bro får inte tågets elektronik slås ut.



Figur 2.5. Det händer att blixten slår ner i flygplan och då måste flygelektroniken tåla de kraftiga störningar som uppstår.

Ett annat viktig problemområde är naturligtvis antenner. En antenn har till uppgift att ta emot och sända ut elektromagnetiska vågor. Antenner ser olika ut beroende på vilka egenskaper de skall ha. En mobiltelefonantenn får exempelvis inte vara allt för komplicerad att tillverka eftersom de skall produceras i stora mängder medan satellitantenner tillverkas i små serier med mycket höga krav på prestanda och låg vikt.



Figur 2.6. Mobiltelefonanvändare är i dubbel bemärkelse i vägen för telefonens radiosignaler. Dels kan sändarkvaliteten försämrans av att huvudet absorberar radiovågor, dessutom måste absorptionen hållas nere för att begränsa eventuella hälsorisker för användaren.



Figur 2.7. Man kan bygga mindre och lättare TV-satelliter om man kan buckla parabolantennerna så att sändningarna bara når de områden man vill täcka.



Figur 2.8. I en vågledarantenn förs antenssignalen in genom en kabel och utbreder sig i vågledaren. Genom att skära upp slit-sar på lämpliga ställen kan man få antenssignalen att stråla ut till omgivningen på ett bra sätt.



Figur 2.9. Antenner på flygplan måste placeras med omsorg så att de fungerar optimalt. De får exempelvis inte störa varandra.

Det finns en mängd olika tillämpningsområden för elektromagnetisk analys förutom EMC och antenner. Exempel på andra tillämpningsområden är mikrovågsugnar, inversa problem, smygflygplan och mikrovågskomponenter.




Figur 2.10. Mikrovågsugnar måste byggas så att man får en jämn fördelning av effekten i ugnen. Tinar man exempelvis mat vill man inte att ena halvan av maten blir överhettad medan andra halvan fortfarande är frusen.



Figur 2.11. I många situationer känner man till vilka geometriska objekt som påverkar lösningen, och det är själva fältet man söker. Det finns dock en klass av problem där man känner lösningen, men inte geometrin. Dessa kallas inversa problem och ett exempel är när man mäter upp fältet ovanför en kabel, för att räkna ut var den ligger nergrävd.



Figur 2.12. Smygflygplan är designade för att vara svårupptäckta med radar. Att bygga flygplan med denna egenskap vore omöjligt utan en god förståelse av elektromagnetisk vågutbredning.



Figur 2.13. Denna mikrovågskomponent fördelar inkommande mikrovågor till de olika grenarna. Utformningen avgör effektfördelningen.

En förutsättning för att den elektromagnetiska utvecklingen skall kunna fortsätta är att vi kan lösa svårare och mer komplexa elektromagnetiska problem. För att studera elektromagnetiska frågeställningar finns det olika hjälpmedel att tillgå. Den hittills mest betydelsefulla är experiment där man genom mätningar kan förstå och analysera de problem som studeras. Men behovet av prototyper för experiment, och det faktum att man inte alltid kommer åt att mäta, har gjort att datorsimule-

ringar i allt större utsträckning kommit att användas. Om man exempelvis redan på konstruktionsstadiet kan beräkna en antens prestanda slipper man kostsamma mätningar och tidsödande omkonstruktioner. Simuleringar kräver dock i allmänhet att man har tillgång till goda datorresurser och avancerade beräkningsprogram.

Datorresurserna har sedan ett halvt sekel förbättrats i en oerhörd takt och det finns ingen anledning att tro att utvecklingen kommer upphöra. Men kraftfullare datorer räcker inte för att kunna göra mer sofistikerade simuleringar. Beräkningsprogrammen måste också förbättras och det är i detta sammanhang min avhandling bidrar till den elektromagnetiska utvecklingen.

Denna avhandling beskriver metoder för att lösa Maxwells ekvationer numeriskt. Ordet *numeriskt* kan tolkas som att fysiken omformuleras för att kunna programmeras och lösas med hjälp av datorer. De fysikaliska lagar Maxwell formulerade är beskrivna med matematiska begrepp, men för att kunna lösa dessa matematiska ekvationer konstruerar vi numeriska metoder. En mycket fundamental egenskap hos numeriska metoder är att de är konsistenta, vilket innebär att numeriken verkligen beskriver samma fenomen som matematiken.

De numeriska metoder som behandlas i denna avhandling har det gemensamt att man delar upp beräkningsrummet i små celler, exempelvis små kuber. I varje cell förenklar man den matematiska formuleringen tillräckligt mycket för att en dator skall kunna tolka problemet. Denna process kallas att diskretisera problemet. En viktig egenskap är att den diskreta lösningen blir bättre och bättre ju finare man diskretiserar.

Med en konvergent numerisk metod skulle man kunna lösa ett problem med i princip hur litet fel som helst under förutsättning att man diskretiserade mycket fint. Men ju finare celler man använder desto fler blir de, om man har ett givet objekt att studera. Med fler celler ökar minnesbehovet samt kraven på datorns beräkningshastighet. Denna konflikt mellan önskan att ha fina celler för att minska felen och önskan att använda få celler för att datorkraven skall vara måttliga kräver en noggrant val diskretisering.

I ovanstående konflikt finner vi motiveringen till den numeriska analysen som akademiskt ämne. Kan vi utveckla gamla metoder och hitta på nya metoder som minskar diskretiseringsfelen vore mycket vunnet under förutsättning att metoderna inte ökar behovet av datorresurser allt för mycket. Om behovet av datorresurser dessutom minskar med nya metoder kan vi kosta på oss finare diskretiseringar och få mycket mindre fel än tidigare.

Denna avhandling innehåller nio kapitel. Första kapitlet och detta kapitel ger en introduktion till elektromagnetiska beräkningar och i tredje kapitlet beskrivs Maxwells ekvationer som alltså utgör den matematiska grunden för avhandlingen. Maxwells ekvationer beskriver hur elektriska och magnetiska fälten är relaterade till varandra i rum och tid. I princip säger ekvationerna två saker, elektriska fältets förändring i tiden bestäms av magnetiska fältets variation i rummet, samt omvänt, magnetiska fältets förändring i tiden ges av elektriska fältets variation i rummet.

I kapitel fyra ges en delvis detaljerad beskrivning av FD-TD metoden (Finita Differenser i TidsDomänen) som är en standardmetod för att lösa Maxwells ekvatio-

ner numeriskt. Metoden går ut på att ersätta de matematiska relationerna mellan förändringar i tiden och variationer i rummet hos det elektromagnetiska fältet med differenser.

Diskretiseringen består i att ersätta fälten, som ju finns överallt och vid alla tidpunkter, med fält som bara definierats i skärningspunkterna i ett ruttmönster, samt vid vissa jämnt fördelade tidpunkter. Man kan alltså ersätta förändringar eller variationer med differenser enligt följande beskrivning. Antag att vi känner till både det elektriska och magnetiska fältet vid en viss tidpunkt. Då kan vi beräkna variationen i rummet genom att samla ihop de angränsande cellernas värden och ta skillnaden mellan dem. Nu har vi fått variationen i rummet hos fälten och det är dags att ta ett tidssteg. Förändringen i tiden kan approximeras med skillnaden mellan ett nytt och ett gammalt värde. Eftersom variationen i rummet redan är beräknat, och det gamla värdet är känt, är det enkelt att få det nya värdet. Matematiskt kan vi skriva detta som $\text{nya värdet} = \text{gamla värdet} + \text{rumsvariationen} \cdot \text{tidssteget}$. Även om denna beskrivning är kraftigt förenklad beskriver den principen för FD-TD metoden. När vi nu vet att variationerna/förändringarna ersätts av differenser (små celler o korta tidssteg krävs) och att man stegar sig fram i tiden blir begreppet Finita Differenser i Tidsdomänen begripligt.

Kapitel fem beskriver en stor FD-TD simulering där en miljard celler använts. En miljard celler är oerhört många celler. Med en miljard gatstenar skulle man kunna stensätta hela Södermalm och Gamla stan. Skulle man istället lägga alla gatstenar på rad skulle de räkna två och ett halvt varv runt jorden.

FD-TD metoden är en mycket effektiv metod men har en allvarlig brist. Eftersom cellerna är kubiska kan man inte diskretisera krökta objekt på ett bra sätt. Man är hänvisad till LEGO-approximationer av geometrier, vilket ibland introducerar stora fel. En lösning på detta problem ges i kapitel sex och sju där hybridmetoder presenteras. Hybridmetoderna utnyttjar ostrukturerade nät med celler av trianglar eller tetraedrar. Dessa ger inte upphov till LEGO-approximationer utan kan med fördel användas för att diskretisera geometrier. De ostrukturerade metodernas stora nackdel är att de är avsevärt mindre effektiva än FD-TD metoden. Hybridmetoderna kombinerar den noggrannare geometridiskretiseringen hos ostrukturerade metoder med effektiviteten hos FD-TD metoden för celler omkring och mellan geometrier.

Kapitel åtta beskriver en ny metod för att räkna på problem med tunna kablar med FD-TD metoden på ett noggrant sätt. Problem där tunna kablar förekommer är exempelvis EMC-problem där man vill undvika att elektromagnetiska störningar inducerar skadliga strömmar i eventuella kablar. Tidigare metoder har krävt att kablarna följt cellkanterna vilket inneburit att diskretiseringsfelen blivit stora för kablar som löpt snett genom beräkningsområdet. Med den nya metoden tillåts kablar vara godtyckligt orienterade

Slutligen, i kapitel nio diskuteras hur man kan visualisera elektromagnetiska simuleringsresultat. En samling färgbilder illustrerar några visualiseringsmetoder och visar samtidigt en del av de geometrier och frågeställningar som avhandlingen behandlar.

Chapter 3

The Maxwell Equations

3.1 The equations

This thesis deals with numerical approximations of electromagnetic phenomena. These are described by the Maxwell equations, see for instance page 323 in Cheng [18],

$$\begin{aligned}\nabla \cdot \mathbf{D} &= \rho, & (\text{Gauss' law}) \\ \nabla \cdot \mathbf{B} &= 0, & (\text{Gauss' law}) \\ \frac{\partial \mathbf{B}}{\partial t} &= -\nabla \times \mathbf{E}, & (\text{Faraday's law}) \\ \frac{\partial \mathbf{D}}{\partial t} &= \nabla \times \mathbf{H} - \mathbf{J}_e, & (\text{Ampere's law})\end{aligned}\tag{3.1}$$

where $\mathbf{E}(\mathbf{x}, t)$ is the electric field [V/m], $\mathbf{D}(\mathbf{x}, t)$ is the electric flux density [C/m²], $\mathbf{H}(\mathbf{x}, t)$ is the magnetic field [A/m], $\mathbf{B}(\mathbf{x}, t)$ is the magnetic flux density [Wb/m²], $\mathbf{J}_e(\mathbf{x}, t)$ is the electric current density [A/m²] and $\rho(\mathbf{x}, t)$ is the charge density [C/m³]. The Maxwell equations are complemented by the equation of continuity,

$$\frac{\partial \rho}{\partial t} = -\nabla \cdot \mathbf{J}_e.\tag{3.2}$$

The two Gauss' laws can be derived from Ampère's law, Faraday's law and the equation of continuity by taking the divergence on Ampère's and Faraday's laws.

For linear, isotropic and non-dispersive materials we have

$$\mathbf{B} = \mu \mathbf{H} \quad \text{and} \quad \mathbf{D} = \epsilon \mathbf{E}.\tag{3.3}$$

Furthermore we allow for materials with isotropic, non-dispersive electric losses that attenuate \mathbf{E} fields via conversion to heat energy. This yields

$$\mathbf{J}_e = \sigma \mathbf{E}.\tag{3.4}$$

Materials for which $\sigma = 0$ are referred to as lossless. Inserting these three relations in (3.1) yields

$$\begin{aligned}
 \nabla \cdot (\epsilon \mathbf{E}) &= \rho, & (\text{Gauss' law}) \\
 \nabla \cdot (\mu \mathbf{H}) &= 0, & (\text{Gauss' law}) \\
 \mu \frac{\partial \mathbf{H}}{\partial t} &= -\nabla \times \mathbf{E}, & (\text{Faraday's law}) \\
 \epsilon \frac{\partial \mathbf{E}}{\partial t} &= \nabla \times \mathbf{H} - \sigma \mathbf{E}, & (\text{Ampere's law})
 \end{aligned} \tag{3.5}$$

where $\mathbf{E} = (E_x, E_y, E_z)$ is the electric field [V/m], $\mathbf{H} = (H_x, H_y, H_z)$ is the magnetic field [A/m], $\epsilon(\mathbf{x})$ is the electric permittivity [F/m], $\mu(\mathbf{x})$ is the magnetic permeability [H/m] and $\sigma(\mathbf{x})$ is the electric conductivity [S/m]. Writing them component by component, we get

$$\left\{ \begin{array}{l}
 \epsilon \frac{\partial E_x}{\partial t} = \frac{\partial H_z}{\partial y} - \frac{\partial H_y}{\partial z} - \sigma E_x, \\
 \epsilon \frac{\partial E_y}{\partial t} = \frac{\partial H_x}{\partial z} - \frac{\partial H_z}{\partial x} - \sigma E_y, \\
 \epsilon \frac{\partial E_z}{\partial t} = \frac{\partial H_y}{\partial x} - \frac{\partial H_x}{\partial y} - \sigma E_z, \\
 \mu \frac{\partial H_x}{\partial t} = \frac{\partial E_y}{\partial z} - \frac{\partial E_z}{\partial y} - \sigma^* H_x, \\
 \mu \frac{\partial H_y}{\partial t} = \frac{\partial E_z}{\partial x} - \frac{\partial E_x}{\partial z} - \sigma^* H_y, \\
 \mu \frac{\partial H_z}{\partial t} = \frac{\partial E_x}{\partial y} - \frac{\partial E_y}{\partial x} - \sigma^* H_z.
 \end{array} \right. \tag{3.6}$$

We have now introduced the equivalent magnetic loss $\sigma^*(\mathbf{x})$ [Ω/m], see Chapter 3 in Taflove [98]. This increases the symmetry of the Maxwell equations though it is not compatible with Gauss' law for the magnetic flux density. We introduce it because our implementation of FD-TD has the capability to include this term.

Yet another way to write (3.6) is,

$$\mathbf{u}_t = A\mathbf{u}_x + B\mathbf{u}_y + C\mathbf{u}_z, \tag{3.7}$$

where $\mathbf{u} = (E_x \ E_y \ E_z \ H_x \ H_y \ H_z)^T$. All matrices $\xi_1 A + \xi_2 B + \xi_3 C$ for any vector ξ with $\xi_1^2 + \xi_2^2 + \xi_3^2 = 1$ have the same six eigenvalues. They are $-c$, $-c$, 0 , 0 , c and c where $c = 1/\sqrt{\mu\epsilon}$ is the speed of propagation for the electromagnetic wave. This means that we need exactly two boundary conditions at any given boundary.

The Maxwell equations is a hyperbolic system because all eigenvalues are real. See [43] for the definition of hyperbolic.

3.2 Reduction to two dimensions

In two dimensions, (3.6) reduces to two independent set of equations, usually referred to as the transverse magnetic (TM) mode and the transverse electric (TE) mode. If we assume that there are no variations in the z-direction, we get the TM_Z mode

$$\begin{cases} \mu \frac{\partial H_x}{\partial t} &= -\frac{\partial E_z}{\partial y} - \sigma^* H_x, \\ \mu \frac{\partial H_y}{\partial t} &= \frac{\partial E_z}{\partial x} - \sigma^* H_y, \\ \epsilon \frac{\partial E_z}{\partial t} &= \frac{\partial H_y}{\partial x} - \frac{\partial H_x}{\partial y} - \sigma E_z, \end{cases} \quad (3.8)$$

and the TE_Z mode

$$\begin{cases} \epsilon \frac{\partial E_x}{\partial t} &= \frac{\partial H_z}{\partial y} - \sigma E_x, \\ \epsilon \frac{\partial E_y}{\partial t} &= -\frac{\partial H_z}{\partial x} - \sigma E_y, \\ \mu \frac{\partial H_z}{\partial t} &= \frac{\partial E_y}{\partial x} - \frac{\partial E_x}{\partial y} - \sigma^* H_z. \end{cases} \quad (3.9)$$

These two modes are decoupled, i.e. they contain no common field component. They are completely independent for isotropic materials, and they can exist simultaneously with no mutual interaction.

3.3 Reduction to one dimension

If we further assume that the magnetic field in (3.8) has no variation in the y-direction, we get

$$\begin{cases} \mu \frac{\partial H_y}{\partial t} &= \frac{\partial E_z}{\partial x} - \sigma^* H_y, \\ \epsilon \frac{\partial E_z}{\partial t} &= \frac{\partial H_y}{\partial x} - \sigma E_z. \end{cases} \quad (3.10)$$

Similar formulas can be derived for other combinations of the fields.

3.4 Integral formulation

The Maxwell equations in (3.5) are given in a partial differential equation (PDE) formulation. It is also possible to cast them in an integral formulation. It can be derived from the PDE formulation: The two Gauss' laws are integrated over an arbitrary fixed control volume after which the divergence theorem is applied to these integrals. Faraday's and Ampère's laws are integrated over a control surface,

S , after which the Stokes theorem is applied to the integrals containing the curl operator. We get

$$\begin{aligned}
\oiint_S \epsilon \mathbf{E} \cdot d\hat{S} &= 0, & \text{(Gauss' law)} \\
\oiint_S \mu \mathbf{H} \cdot d\hat{S} &= 0, & \text{(Gauss' law)} \\
\frac{\partial}{\partial t} \iint_S \mu \mathbf{H} \cdot d\hat{S} &= - \oint_C \mathbf{E} \cdot d\hat{l} - \iint_S \sigma^* \mathbf{H} \cdot d\hat{S}, & \text{(Faraday's law)} \\
\frac{\partial}{\partial t} \iint_S \epsilon \mathbf{E} \cdot d\hat{S} &= \oint_C \mathbf{H} \cdot d\hat{l} - \iint_S \sigma \mathbf{E} \cdot d\hat{S}, & \text{(Ampere's law)}
\end{aligned} \tag{3.11}$$

where C is the contour that bounds the surface S . The surface S in the Gauss' laws is not the same as the S in the Faraday's and Ampère's laws. In Gauss' laws, it is the surface of the control volume. This integral formulation of the Maxwell equations is used to construct several of the numerical methods treated in this thesis. It is possible to get other integral formulations of the Ampère's and Faraday's laws, for instance by integrating them over a volume instead of a surface.

Formulas (3.8), (3.9) and (3.10) can also be cast in integral formulations in a similar manner.

3.5 The wave equation

If we take the time derivative of Ampère's law in (3.5) and assume that the material properties are time independent, we obtain

$$\epsilon \frac{\partial^2 \mathbf{E}}{\partial t^2} = -\nabla \times \frac{1}{\mu} \nabla \times \mathbf{E} - \sigma \frac{\partial \mathbf{E}}{\partial t}. \tag{3.12}$$

For lossless homogeneous materials without sources this reduces to

$$\frac{\partial^2 \mathbf{E}}{\partial t^2} = c^2 \Delta \mathbf{E}, \tag{3.13}$$

where $c = 1/\sqrt{\mu\epsilon}$ is the speed of propagation for the electromagnetic wave. In a similar manner, we may show that

$$\frac{\partial^2 \mathbf{H}}{\partial t^2} = c^2 \Delta \mathbf{H} \tag{3.14}$$

for lossless homogeneous materials without sources.

3.6 Material properties

In (3.6) we have four parameters. For vacuum they are $\mu \equiv \mu_0 = 4\pi \cdot 10^{-7} \text{ Vs/Am}$, $\epsilon \equiv \epsilon_0 \approx 10^{-9}/36\pi \approx 8.8541878 \cdot 10^{-12} \text{ As/Vm}$, $\sigma^* = 0 \text{ }\Omega/\text{m}$ and $\sigma = 0 \text{ S/m}$. The speed of light in vacuum is defined by $c_0 = 2.99792458 \cdot 10^8 \text{ m/s} \approx 1/\sqrt{\mu_0\epsilon_0}$. For other materials it is customary to define their permeability and permittivity relative to those of vacuum, i.e. we have $\epsilon = \epsilon_r\epsilon_0$ and $\mu = \mu_r\mu_0$. The relative permittivities for some common materials are listed in Table 3.1. The data have been taken from Table B-3 in Cheng [18]. For most materials, ϵ_r and μ_r are frequency dependent. Materials for which we assume that ϵ_r and μ_r are independent of frequency are referred to as simple materials. Frequency-dependent materials will be briefly addressed in Chapter 4.12. The values listed in Table 3.1 are average low-frequency values at room temperature. Note that we always have $\epsilon_r \geq 1$. Most materials where $\mu_r \neq 1$ are metals with high conductivity. We treat these materials as perfect electric conductors.

Material	ϵ_r
Teflon	2.1
Rubber	2.3-4.0
Bakelite	5.0
Distilled Water	80

Table 3.1. Relative permittivities for some common materials.

At the interface between two lossless media (we have, see Table 7-3 on page 330 in Cheng [18])

$$\begin{aligned}
 \mathbf{n} \cdot (\mathbf{D}_1 - \mathbf{D}_2) &= 0, \\
 \mathbf{n} \times (\mathbf{E}_1 - \mathbf{E}_2) &= 0, \\
 \mathbf{n} \cdot (\mathbf{B}_1 - \mathbf{B}_2) &= 0, \\
 \mathbf{n} \times (\mathbf{H}_1 - \mathbf{H}_2) &= 0,
 \end{aligned} \tag{3.15}$$

where the subscripts indicate which region the field belongs to, and \mathbf{n} is the interface normal. Using the relations in (3.3), we get

$$\begin{aligned}
 \mathbf{n} \cdot (\epsilon_1 \mathbf{E}_1 - \epsilon_2 \mathbf{E}_2) &= 0, \\
 \mathbf{n} \times (\mathbf{E}_1 - \mathbf{E}_2) &= 0, \\
 \mathbf{n} \cdot (\mu_1 \mathbf{H}_1 - \mu_2 \mathbf{H}_2) &= 0, \\
 \mathbf{n} \times (\mathbf{H}_1 - \mathbf{H}_2) &= 0.
 \end{aligned} \tag{3.16}$$

For perfect electric conductors (PEC) we have (compare with Table 7-4 on page 331 in [18])

$$\begin{aligned}
 \mathbf{n} \cdot \epsilon \mathbf{E} &= \rho_s , \\
 \mathbf{n} \times \mathbf{E} &= 0 , \\
 \mathbf{n} \cdot \mathbf{H} &= 0 , \\
 \mathbf{n} \times \mathbf{H} &= \mathbf{J}_s ,
 \end{aligned}
 \tag{3.17}$$

where ρ_s is the surface charge density [C/m^2] and \mathbf{J}_s is the surface current density [A/m]. Note that the normal \mathbf{n} is pointing out from the PEC region. It may seem odd that we have six boundary conditions when we should only have two. However the first and fourth conditions are not true boundary conditions, because ρ_s and \mathbf{J}_s are unknown, and the third condition can easily be shown to be a consequence of the second condition.

PECs are characterized by having no tangential electric field at the surface. This is a consequence of the term perfect conductor. If there were a tangential electric field it would drive an infinite surface current which is clearly unphysical. However, this does not imply that the surface current must be zero. In fact, if there is an external field there will always be surface currents since the magnetic field does only have tangential components at the PEC surface and the surface current is related to the tangential magnetic field through the fourth condition in (3.17).

Chapter 4

FD-TD

This Chapter was written prior to the publication of the second edition of Taflove's book on FD-TD [100]. Hence all references are to the first edition [98].

4.1 Introduction to FD-TD

The most commonly used time-domain method for solving the Maxwell equations is the Finite-Difference Time-Domain (FD-TD) method. It was introduced by Yee in 1966 [110] and is sometimes referred to as the Yee scheme. The method was further developed and promoted by Taflove in the 1970s, and he also coined the acronym FD-TD.

Several books have been published dealing with the FD-TD scheme [56, 98, 99, 52, 97]. The survey paper by Shlager and Schneider that appeared in [99] illustrates the rapid growth in the use of FD-TD.

The FD-TD method has been attractive for industrial users since the early 1980s because the basic method is relatively simple to program and because the geometry handling is fairly straightforward. The method can also be efficiently implemented on vector computers which made it feasible to solve complex problems on the early supercomputers. As an example, in 1987 SAAB performed lightning analysis on the Swedish fighter aircraft Gripen on a grid with approximately $60 \times 30 \times 30$ cells.

4.2 Discretization used in FD-TD

The FD-TD scheme is an explicit finite difference scheme using central differences on a staggered Cartesian grid (both space and time), i.e. it is a leap-frog scheme. It is second-order accurate in both time and space. "Staggered" here indicates that the different electromagnetic components are not located at the same place (see Figures 4.1 and 4.2). Furthermore, the fields are not represented on the same time levels (see Figure 6.1).




Figure 4.1. The FD-TD TM grid, with problem size $N_x = 8$ and $N_y = 5$. The Huygens' surfaces (dashed line) are placed in the second cell from the outer boundary.

Staggering the variables in the computational grid is a straightforward consequence of the nature of the Maxwell equations and the central finite differences. If the leap-frog scheme is applied on a grid that is not staggered, it would result in $2 \cdot 8$ uncoupled discrete equations. Time and memory are saved by only solving one of these. (memory savings: 8, time step savings: 2 per unknown $\Rightarrow 16$)

The main drawback of the FD-TD scheme is the inability to represent curved boundaries and small geometrical details. Curved objects must be modeled by staircasing, i.e. they must fit into the Cartesian grid and hence look like they were made of Lego blocks. As will be demonstrated in Chapter 6, staircasing destroys the second-order accuracy.

Figure 4.1 shows a grid for the 2D TM equations. The dashed line indicates the limit between the total field region and the scattered field region. Further explanation can be found in Section 4.7.

One cell of a 3D FD-TD grid is given in Figure 4.2. The magnetic field components are defined on the cell's faces, and the electric field components are defined on the cell's edges. This choice is arbitrary and does not affect the behavior of the scheme. For a computational domain with the number of cells (N_x, N_y, N_z) , we




Figure 4.2. Positions of the electric and magnetic field vector components in a unit Yee cell.

define

$$\begin{aligned}
E_x|_{i+\frac{1}{2},j,k}^n, & \quad i = 1, \dots, N_x, \quad j = 1, \dots, N_y + 1, \quad k = 1, \dots, N_z + 1, \\
E_y|_{i,j+\frac{1}{2},k}^n, & \quad i = 1, \dots, N_x + 1, \quad j = 1, \dots, N_y, \quad k = 1, \dots, N_z + 1, \\
E_z|_{i,j,k+\frac{1}{2}}^n, & \quad i = 1, \dots, N_x + 1, \quad j = 1, \dots, N_y + 1, \quad k = 1, \dots, N_z, \\
H_x|_{i,j+\frac{1}{2},k+\frac{1}{2}}^{n-\frac{1}{2}}, & \quad i = 1, \dots, N_x + 1, \quad j = 1, \dots, N_y, \quad k = 1, \dots, N_z, \\
H_y|_{i+\frac{1}{2},j,k+\frac{1}{2}}^{n-\frac{1}{2}}, & \quad i = 1, \dots, N_x, \quad j = 1, \dots, N_y + 1, \quad k = 1, \dots, N_z, \\
H_z|_{i+\frac{1}{2},j+\frac{1}{2},k}^{n-\frac{1}{2}}, & \quad i = 1, \dots, N_x, \quad j = 1, \dots, N_y, \quad k = 1, \dots, N_z + 1,
\end{aligned} \tag{4.1}$$

where $n = 0, \dots, N_t$ for all six components and N_t is the number of time steps taken. Initial values are needed for $\mathbf{H}^{-\frac{1}{2}}$ and \mathbf{E}^0 . Note that our notation is slightly different from that used in [98]. In our notation there is always a direct correspondence between the indexes and the physical location of a field component. For example,

$$H_x|_{i,j+\frac{1}{2},k+\frac{1}{2}}^{n-\frac{1}{2}} \text{ is located at } ((i-1)\Delta x, (j-1/2)\Delta y, (k-1/2)\Delta z), \tag{4.2}$$

at $t = (n-1/2)\Delta t$ where Δx , Δy and Δz are the spatial cell sizes and Δt is the time increment. The total spatial problem size is $N = N_x N_y N_z$. The storage space needed for this is approximately $24N$ byte for 32-bit precision and $48N$ byte for 64-bit precision.

4.3 The leap-frog scheme

In homogeneous materials with $\sigma = \sigma^* = 0$, the following formulas comprise the FD-TD updating stencils for the electromagnetic field components.

$$\begin{aligned} H_x|_{i,j+\frac{1}{2},k+\frac{1}{2}}^{n+\frac{1}{2}} &= H_x|_{i,j+\frac{1}{2},k+\frac{1}{2}}^{n-\frac{1}{2}} + \frac{\Delta t}{\mu\Delta z} \left[E_y|_{i,j+\frac{1}{2},k+1}^n - E_y|_{i,j+\frac{1}{2},k}^n \right] \\ &\quad - \frac{\Delta t}{\mu\Delta y} \left[E_z|_{i,j+1,k+\frac{1}{2}}^n - E_z|_{i,j,k+\frac{1}{2}}^n \right] \end{aligned} \quad (4.3)$$

$$\begin{aligned} H_y|_{i+\frac{1}{2},j,k+\frac{1}{2}}^{n+\frac{1}{2}} &= H_y|_{i+\frac{1}{2},j,k+\frac{1}{2}}^{n-\frac{1}{2}} + \frac{\Delta t}{\mu\Delta x} \left[E_z|_{i+1,j,k+\frac{1}{2}}^n - E_z|_{i,j,k+\frac{1}{2}}^n \right] \\ &\quad - \frac{\Delta t}{\mu\Delta z} \left[E_x|_{i+\frac{1}{2},j,k+1}^n - E_x|_{i+\frac{1}{2},j,k}^n \right] \end{aligned} \quad (4.4)$$

$$\begin{aligned} H_z|_{i+\frac{1}{2},j+\frac{1}{2},k}^{n+\frac{1}{2}} &= H_z|_{i+\frac{1}{2},j+\frac{1}{2},k}^{n-\frac{1}{2}} + \frac{\Delta t}{\mu\Delta y} \left[E_x|_{i+\frac{1}{2},j+1,k}^n - E_x|_{i+\frac{1}{2},j,k}^n \right] \\ &\quad - \frac{\Delta t}{\mu\Delta x} \left[E_y|_{i+1,j+\frac{1}{2},k}^n - E_y|_{i,j+\frac{1}{2},k}^n \right] \end{aligned} \quad (4.5)$$

$$\begin{aligned} E_x|_{i+\frac{1}{2},j,k}^{n+1} &= E_x|_{i+\frac{1}{2},j,k}^n - \frac{\Delta t}{\epsilon\Delta z} \left[H_y|_{i+\frac{1}{2},j,k+\frac{1}{2}}^{n+\frac{1}{2}} - H_y|_{i+\frac{1}{2},j,k-\frac{1}{2}}^{n+\frac{1}{2}} \right] \\ &\quad + \frac{\Delta t}{\epsilon\Delta y} \left[H_z|_{i+\frac{1}{2},j+\frac{1}{2},k}^{n+\frac{1}{2}} - H_z|_{i+\frac{1}{2},j-\frac{1}{2},k}^{n+\frac{1}{2}} \right] \end{aligned} \quad (4.6)$$

$$\begin{aligned} E_y|_{i,j+\frac{1}{2},k}^{n+1} &= E_y|_{i,j+\frac{1}{2},k}^n - \frac{\Delta t}{\epsilon\Delta x} \left[H_z|_{i+\frac{1}{2},j+\frac{1}{2},k}^{n+\frac{1}{2}} - H_z|_{i-\frac{1}{2},j+\frac{1}{2},k}^{n+\frac{1}{2}} \right] \\ &\quad + \frac{\Delta t}{\epsilon\Delta z} \left[H_x|_{i,j+\frac{1}{2},k+\frac{1}{2}}^{n+\frac{1}{2}} - H_x|_{i,j+\frac{1}{2},k-\frac{1}{2}}^{n+\frac{1}{2}} \right] \end{aligned} \quad (4.7)$$

$$\begin{aligned} E_z|_{i,j,k+\frac{1}{2}}^{n+1} &= E_z|_{i,j,k+\frac{1}{2}}^n - \frac{\Delta t}{\epsilon\Delta y} \left[H_x|_{i,j+\frac{1}{2},k+\frac{1}{2}}^{n+\frac{1}{2}} - H_x|_{i,j-\frac{1}{2},k+\frac{1}{2}}^{n+\frac{1}{2}} \right] \\ &\quad + \frac{\Delta t}{\epsilon\Delta x} \left[H_y|_{i+\frac{1}{2},j,k+\frac{1}{2}}^{n+\frac{1}{2}} - H_y|_{i-\frac{1}{2},j,k+\frac{1}{2}}^{n+\frac{1}{2}} \right] \end{aligned} \quad (4.8)$$

For lossy materials the electric field update equations are modified, for example (4.6), is replaced by

$$\begin{aligned} E_x|_{i+\frac{1}{2},j,k}^{n+1} &= \frac{1 - \frac{\sigma\Delta t}{2\epsilon}}{1 + \frac{\sigma\Delta t}{2\epsilon}} E_x|_{i+\frac{1}{2},j,k}^n - \frac{\Delta t/\epsilon\Delta z}{1 + \frac{\sigma\Delta t}{2\epsilon}} \left[H_y|_{i+\frac{1}{2},j,k+\frac{1}{2}}^{n+\frac{1}{2}} - H_y|_{i+\frac{1}{2},j,k-\frac{1}{2}}^{n+\frac{1}{2}} \right] \\ &\quad + \frac{\Delta t/\epsilon\Delta y}{1 + \frac{\sigma\Delta t}{2\epsilon}} \left[H_z|_{i+\frac{1}{2},j+\frac{1}{2},k}^{n+\frac{1}{2}} - H_z|_{i+\frac{1}{2},j-\frac{1}{2},k}^{n+\frac{1}{2}} \right], \end{aligned} \quad (4.9)$$

where the conductivity term σE_x is discretized using the average of the electric field at time levels n and $n+1$. Using only the value from time level n gives an

unstable scheme, and using only the value from time level $n + 1$ gives a noncentered scheme. For highly lossy media one could take into account the rapid exponential decrease of field strengths and introduce a scaling of the variables; this would yield the so-called exponential timestepping scheme. However, this scheme does not give any significant improvements [80] and is thus not further discussed.

Comparing (4.6) and (4.9) we see that one extra arithmetic operation is needed for lossy material. For inhomogeneous materials, we replace ϵ in (4.9) with $\epsilon_{i+\frac{1}{2},j,k}$ and similarly for σ . Note that we need an ϵ -value for each of the three electric field component updates in a cell. These three ϵ -values will differ in the vicinity of a material interface. Chapter 9 in the Ph. D. thesis of Ulf Andersson [5] gives a detailed analysis of how to calculate these discrete values.

4.4 Stability conditions

Because FD-TD is an explicit scheme, there is a limit on the time step Δt to ensure stability. It is given by:

$$\Delta t < \frac{1}{c\sqrt{\frac{1}{(\Delta x)^2} + \frac{1}{(\Delta y)^2} + \frac{1}{(\Delta z)^2}}}, \quad (4.10)$$

where c is the wave propagation speed. We define the CFL number as

$$CFL = c\Delta t\sqrt{\frac{1}{(\Delta x)^2} + \frac{1}{(\Delta y)^2} + \frac{1}{(\Delta z)^2}}, \quad (4.11)$$

and may thus write the stability condition as

$$CFL < 1. \quad (4.12)$$

CFL stands for *Courant-Friedrichs-Lewy* (see page 54 in [43]).

4.5 Performance of the leap-frog update

The leap-frog update is the core of an FD-TD solver, and therefore it must be implemented as efficiently as possible. There are two major obstacles in getting an efficient implementation. The updating stencils of (4.3)–(4.8) consist of two multiplications and four additions/subtractions each. Most computers today are constructed to perform the same number of multiplications and additions in every clock cycle. In our case, this means that we can achieve at most 75% of the peak performance. The other main obstacle is the need for memory bandwidth. For instance, to compute (4.3) we need to fetch five field values from memory and store one field value to memory. Most computers cannot do this as quickly as they perform the calculations. The constant coefficients will reside in registers and need not be fetched from memory for every update.

It is possible to reduce the number of multiplications in (4.3) by one by scaling the fields with the cell size. Let $\tilde{E}_z = \Delta z E_z$, etc. We get

$$\tilde{H}_x|_{i,j+\frac{1}{2},k+\frac{1}{2}}^{n+\frac{1}{2}} = \tilde{H}_x|_{i,j+\frac{1}{2},k+\frac{1}{2}}^{n-\frac{1}{2}} + \frac{\Delta t}{\mu} \left[\begin{array}{l} \tilde{E}_y|_{i,j+\frac{1}{2},k+1}^n - \tilde{E}_y|_{i,j+\frac{1}{2},k}^n - \\ \tilde{E}_z|_{i,j+1,k+\frac{1}{2}}^n + \tilde{E}_z|_{i,j,k+\frac{1}{2}}^n \end{array} \right]. \quad (4.13)$$

However, this reduction will lead to little or no gain in execution time on most modern computers because the number of additions still dominates. This is illustrated in Table 4.1.

We will now present some illustrating performance results for the leap-frog update. All tests are performed on the same “standard” problem. We set $N_x = N_y = N_z = 100$ and $N_t = 100$. We use a point source for excitation and the Mur first-order ABC [74] as grid terminator. Timing is performed over the timestepping loop, i.e. we omit initialization and post processing. We also omit the first time step from the timing (timing is performed over 99 iterations), because it might contain initialization overhead. All calculations are performed in 64-bit precision.

Table 4.1 shows the effect on execution time of a reduction in the number of multiplications per component update by one. On the IBM processors, there are no gain in execution time. On the Sun there is some gain, but only about 4%. This should be compared to the 17% (compare (4.3) and (4.13)) decrease in the number of floating point operations.

Computer	Reduced code	Original code
IBM pwr3, 200 MHz	23.13 sec.	23.07 sec.
IBM pwr2, 160 MHz	22.21 sec.	22.24 sec.
Sun Ultra 1, 167 MHz	93.28 sec.	97.00 sec.

Table 4.1. Execution times for the leap-frog update for lossless homogeneous material. These test were performed in February 2000.

For lossless materials we have 36 arithmetic operations per cell. For lossy materials we have 42 arithmetic operations per cell if both σ and σ^* are nonzero. Obviously, lossy inhomogeneous materials increase the execution time. This effect is illustrated in Table 4.2. The number of floating point operations per iteration (Flop/iteration) includes the operations performed by the first-order Mur ABC, see Section 4.6.

Material	Lossless homogeneous	Lossy inhomogeneous
Flop/cell	36	42
Flop/iteration	36 000 900	41 941 200
Performance (Mflop/s)	120.01	93.31
Percentage of peak perf.	13.6	10.6
Time (s)	29.70	44.50

Table 4.2. Performance for homogeneous and inhomogeneous materials on an IBM pwr3 222 MHz processor. These test were performed in June 2000 and cannot be compared to the results in Table 4.1 because different processors and compiler versions were used.

4.6 Boundary conditions

Perfect electric conductors (PEC) are characterized by the absence of tangential electric field at the surface, as discussed in Section 3.6. A PEC must be described using a staircase approximation to fit into the FD-TD scheme. This staircase procedure is a major cause of errors in FD-TD calculations.

It is possible to model PECs by changing the coefficients in (4.3)–(4.5), but a more efficient implementation for homogeneous materials is to first update all the electric field components using (4.6)–(4.8) and then set all \mathbf{E} fields on the surface of the object to zero.

Many applications involve geometries with unlimited surroundings. These situations are called open problems. In these cases it is necessary to limit the computational domain by introducing an artificial outer boundary. At this boundary we need to apply a boundary condition, and this condition should be designed to absorb outgoing waves. Hence we refer to it as an absorbing boundary condition (ABC). One could also think of this boundary as having the property of not reflecting any outgoing waves back into the computational domain and hence name it a nonreflecting boundary condition (NRBC).

The history of absorbing boundary conditions for the FD-TD scheme is carefully covered in Chapter 7 of [98]. This chapter concludes with a description of the perfectly matched layer (PML) introduced by Berenger [9] in 1994, which was a tremendous breakthrough in ABC methodology. The basic idea is to surround the computational domain with an absorbing layer. This concept had been tried before, but there were problems with reflections in the interface between the computational domain and the absorbing layer [50]. The key to the success of PML is that there are no reflections at this interface, at least not for the continuous problem. This is true for all frequencies and all angles of incidence.

One of the chapters in [99] covers the further development of PML. The original formulation of PML is a weakly hyperbolic system [1], which might cause stability problems. This formulation is based on splitting the six field components into two parts each. Later formulations instead introduced a lossy anisotropic absorbing material [38]. They are referred to as unsplit PML (U-PML).

The U-PML formulation relies on the fact that it is possible to derive matching conditions for lossy anisotropic (uniaxial) media, so that incident plane waves are purely transmitted. The reflectionless conditions for the permittivity and permeability in the uniaxial media are found to be:

$$\bar{\epsilon} = \bar{\mu} = \begin{bmatrix} a & 0 & 0 \\ 0 & a & 0 \\ 0 & 0 & a^{-1} \end{bmatrix}. \quad (4.14)$$

The parameter a is then chosen to be lossy. A natural choice in the frequency domain is

$$a = 1 + \frac{\sigma}{j\omega\epsilon_0}, \quad (4.15)$$

where ω is the frequency.

The attenuation of the PML depends on the size of the lossy parameter (σ), the depth of the absorbing layer and the angle of incidence of the wave. The value of σ should be as large as possible to improve absorption. However, this would result in a step discontinuity in σ in the transition between the interior region and the absorbing layer. In the discrete space, this leads to large reflections of the fields. The parameter σ is therefore chosen as a smoothly increasing function, starting from zero.

A systematic way to evaluate the performance of different ABCs was given in [72]. Their test consisted of using a point source in 2D and comparing the results with numerical results from a larger domain. We have used this test case on a number of different ABCs. The result is presented in Figure 4.3. All the different ABCs are described in Chapter 7 of [98], except the U-PML scheme which is described in [38]. The Mei Fang result presented in Figure 4.3 has been obtained by applying the Mei Fang procedure to the second-order Mur scheme. The notation U-PML X refers to a U-PML with a layer of X cells. The U-PML results are equivalent to what we would have obtained using PML. It is evident from the results in Figure 4.3 that it is possible to achieve much better absorption with U-PML/PML than with previously developed ABCs. Figure 4.3 is the same type of graph as Figure 5b in [72] and Figure 7.8 in [98]. Note that formula (7.46) in [98] contains a misprint. The factor should be 1/320 not 1/32 (see [72]).

The U-PML can be extended to treat frequency dispersive materials, as shown in Section 5.9 in [99]. With U-PML it is possible to construct an arbitrarily good ABC by increasing the number of cells in the U-PML layer. However, there is an increase in cost when increasing the thickness of the layer.

A perfectly matched layer must be terminated at its outer boundary, and one possibility would be to use a classical nonreflecting boundary condition to terminate the U-PML layer. But this is seldom done because the extra cost of implementing and performing this is higher than simply terminating the outer boundary with a PEC condition. This is not as bad as it first might appear. A wave propagating from the inner of the domain will be attenuated exponentially during propagation through the U-PML, and when the wave is reflected in the outer boundary, it will




Figure 4.3. Comparison of different ABCs for the TM Maxwell equations.

be further attenuated on its way back to the inner domain. If this damping is not enough, we just add another cell layer of U-PML, and the effect is still better than using a classical nonreflecting boundary condition to terminate the U-PML layer.

The Gems time domain 3D code includes PML, U-PML for dispersive materials, and also the first-order Mur ABC [74]. The Mur scheme is the application of the Engquist-Majda [31] ABC to the Maxwell equations.

Other approaches to ABC are still being explored in the search for a cheaper ABC, for example [42, 82]. We have not explored this area. One interesting approach is to use the plane wave time domain (PWTD) method [93] as ABC. This would make it possible to put the ABC only a few cells from the scattering object. On the other hand, this is a global ABC. Traditionally, global ABCs have been considered too computationally expensive. This obstacle may be overcome by PWTD, but we are not quite there yet.

It is well known that it is possible to achieve very good ABCs by using integral formulations. However, if an almost perfect ABC consumes 90% of the available computer resources, it is better to spend some of these resources on a finer discretization.

4.7 Sources

There are many different ways to excite the fields depending on what we want to simulate. Excitations that we use include Huygens' surfaces to model plane waves, point sources to model dipoles, and current and voltage sources in thin wires. Waveguide excitations, which are a major part of the GEMS project, are briefly discussed in Section 7.3.1.

Incident plane waves are generated by Huygens' surfaces. They are carefully described in Chapter 6.5 of [98]. However, he never calls them "Huygens' surfaces". Instead he has a more mathematical viewpoint and refers to it as "Total-Field/Scattered-Field Formulation".

A very simple excitation is to use a point source. It can for instance be used to model a dipole. This is done by adding a source term to one of the electric fields. For example,

$$E_x|_{i_s+\frac{1}{2},j_s,k_s}^{n+1} = E_x|_{i_s+\frac{1}{2},j_s,k_s}^{n+1} - \frac{\Delta t f(t_{n+\frac{1}{2}})}{\epsilon V}, \quad (4.16)$$

where $V = \Delta x \Delta y \Delta z$ is the cell volume. Because the discrete time derivative of E_x is centered around time level $n + 1/2$, we evaluate the source function $f(t)$ at $t = (n + 1/2)\Delta t$. This is necessary in order to retain the second-order convergence.

We distinguish between hard sources and soft sources. A soft point source is characterized by adding a source term to the field equation as in (4.16). A hard point source is characterized by setting the field to a source term and hence overwriting the value given by the leap-frog update.

In order to avoid introducing high frequency components in the numerical solution, it is important to use a smooth source. The source must be zero at $t = 0$ when we start our simulation (unless we combine it with suitably chosen initial values). This introduces a discontinuity at $t = 0$, but this discontinuity can be kept at machine precision level by suitable choices of parameters for the source.

4.8 Visualization as debugging and validation tool

There are several reasons to visualize results produced by electromagnetic computations. We distinguish between three different purposes of visualization. The first purpose is the most obvious which is to present research to other scientists, funding agencies or the general public. The second purpose is perhaps the most thrilling, which is to learn more about the features of the Maxwell equations. This often stems from working with the results in order to present a good visualization. The third purpose is often regarded as unglamorous, but is nevertheless important. Given a certain expectation of the result, visualizations can be used for debugging and validation of the code. This is a very useful technique to find out which part of the code that is incorrect. In this section we will focus on visualization as debugging and validation tool.

There is an intrinsic problem in visualization of FD-TD data, which originates from staggering the electromagnetic components. Saving E_x , E_y and E_z to file for a given time step is straightforward, but doing visualizations based on these staggered components can be misleading. Simply taking a Cartesian slice through the data and plotting one component seldom poses any problem if the half-cell bias is not important. But plotting the magnitude of the electric field requires either some averaging to get the field at the corners of the FD-TD cell or acceptance for the half-cell bias of the components. The latter is equivalent to a onesided interpolation to the cell corner. The first approach blurs the effect of errors in a single component, and the second is more difficult to combine with exact geometric representations. However, bearing this in mind, debugging still benefits from studying the magnitude of staggered components.

A common techniques for debugging and validation is to compare obtained results with a reference solution pointwise, i.e. plot $(u_{\text{code}} - u_{\text{ref}})$ using slice planes or isosurfaces. Analytic solutions are in some cases available but often numerical results are the only option. Numerical results might come from other codes or obtained from the code prior to the feature examined.

Another emerging technology is virtual reality, which provides tools and display systems for an immersive exploration of data. This might be judged as only a hyped technology that is too complicated for anyone but experts to use. The truth is that used correctly, it enhances the possibility of sharing results with others, especially if a complicated geometry is included.

4.9 Parallelization

The core of the FD-TD scheme is relatively straightforward to parallelize. This issue is studied in detail in a separate chapter of the Ph. D. thesis of Ulf Andersson [5], and we give a brief summary here. Chapter 5 in [52] is partly devoted to parallelization of the FD-TD scheme. It also covers vectorization and other optimization issues.

Scale-up and speed-up are different measures of efficiency of a parallel code. We define them by:

Definition 4.1 (Scale-up). *The problem size, N , is increased linearly with the number of processors.*

Definition 4.2 (Speed-up). *The problem size, N , is kept fixed, independent of the number of processors.*

We will discuss parallelization of the leap-frog update. Adding PEC, the first-order Mur ABC, and a point source will not affect our discussion, but adding Huygens' surfaces and PML will.

On a parallel computer, it is easy to achieve perfect scale-up for the leap-frog update. This means that the time to execute a certain number of time steps is constant, when the number of processors (and the problem size) is increased. This

is a rather nice result, because the need for more memory is one major reason for using parallel computers. Many applications are such that the time to complete them is acceptable if we can fit them into memory.

Consider the one-billion-cell computation in Appendix 5. The time to complete this calculation was only slightly more than one hour. Here the number of time steps was chosen so that the wave would sweep past the airplane once. If we were to use this calculation for practical purposes, we would have to increase the number of time steps to allow for reflections, surface waves, etc. to evolve. However, even if we increase the number of time steps by a factor of ten, the calculation could still be performed overnight.

Perfect speed-up is usually not possible to achieve for the Yee scheme. The main reason for this is the low number of arithmetic operations per cell and time step. For homogeneous materials with $\sigma = \sigma^* = 0$, we only perform 36 arithmetic operations per cell and time step. This can be compared to computational fluid dynamics, where a Navier-Stokes solver performs thousands of arithmetic operations per cell and time step.

The parallelization of a full FD-TD solver, including subcell models, near-to-far-field transformation, Huygens' surfaces, etc. is a much more complex problem. It gets even more complicated when the hybrid schemes described in Chapters 6 and 7 are parallelized. The tricky part is how to achieve a good load balancing. This is more or less automatic for the leap-frog update, since the same amount of work is performed in every cell.

4.10 Subcell models

In the numerical simulation of electromagnetic wave propagation, the existence of subgrid scale phenomena poses some difficulties. Subgrid scale phenomena refer to geometrical features that should influence the solution on the computational grid but have length scales shorter than the grid size. For some problems, such as narrow slots, thin material sheets, surface impedances and thin wires there are subcell models developed.

A thin-wire model permitting arbitrarily oriented wires is presented in Chapter 8, but the other subcell models will not be further addressed in this thesis. However, they are thoroughly described in Chapter 10 in Taflov's book [98].

4.11 Near-to-far-field transformations

Three different near-to-far-field transformations have been implemented within the GEMS project. The three transforms are the frequency-domain transform (FD), the time-domain transform (TD) and the continuous-wave transform (CW). These transforms have different applicabilities. The TD transform is suitable when the far field is desired over a range of frequencies, but only for a small number of directions.

For the FD transform, the opposite is true. The CW transform is a special case of the FD transform, in which it is assumed that only one frequency is present in the solution. The use of this transform is somewhat limited, since one of the major reasons for using a time-domain method is the possibility of computing a large number of frequencies simultaneously.

The FD transform is equipped with dispersion compensation. This procedure is described in [63]. It can yield significant improvements, especially for the forward scattering direction. The TD transform is described in [65].

All transforms probe the fields on a surface. The FD transform performs a discrete Fourier transform (DFT) on surface currents enclosing the object during the timestepping. After the timestepping the field is transformed to far field. For the CW transform, the DFT is only performed during the final period of the computation. The TD transform performs a near-to-far-field transformation in the time domain during each time step. After timestepping a fast Fourier transform (FFT) is performed at each far field direction. This gives us the far field for a broad spectrum of frequencies.

4.12 Frequency-dispersive materials

The Yee scheme is constructed from a non-dispersive formulation of the Maxwell equations and hence cannot be used for computation on frequency-dependent materials. However, it is possible to extend the Yee scheme to handle these materials (see Chapter 9 in [99]). This thesis will not treat dispersive materials.

4.13 Divergence-free nature

If $\sigma = \rho = 0$ in (3.5), the divergence of \mathbf{E} and \mathbf{H} should be zero. A very nice property of the FD-TD method is that it preserves the divergence; i.e., if it is zero initially, it will stay zero. A proof of this is given in Section 3.6.9 of [98]. The proof is presented for free space, but it is also valid for inhomogeneous materials. It is actually valid independent of how the discrete ϵ -values are chosen due to cancellation of terms in the proof.

Divergence is a continuous property. The proof is of course applied to a discrete approximation of the divergence. However, this approximation is the most natural way to approximate the divergence.

Chapter 5

Large Scale FD-TD — A Billion Cells

This chapter is a reprint of Paper 6 in the list on Page 10.

Abstract

We describe an FD-TD simulation of a SAAB 2000 aircraft using more than one billion cells. This was performed on 125 nodes of an IBM SP.

We also give a speed-up result for the IBM SP and discuss the possibilities for visualization of FD-TD solutions in a VR-CAVE.

5.1 Introduction

The Swedish center of excellence Parallel and Scientific Computing Institute - (PSCI) [81] conducts a research project within its CEM program called “Large Scale FD-TD”. Within this project a 3D FD-TD code, `pscyee`, has been implemented using Fortran 90.

The capabilities of the sequential version of `pscyee` includes Huygens’ surfaces, PML according to Berenger [9] as absorbing boundary condition (ABC), a generalized thin-wire subcell model (Paper 5 on Page 10) based on the approach of Holland-Simpson [49] and PEC surfaces. Excitation may also be performed with current/voltage sources in the thin-wire model and with point sources. The first-order Mur ABC is available for calculations where the accuracy demand on the ABC is not so high.

Implementation of models for frequency dispersive materials is ongoing. It has already been implemented in a 2D version of the code. `pscyee` will be used as a component in a hybrid time domain solver which will combine the FD-TD scheme

with an implicit Finite-Element Time-Domain (FE-TD) scheme and an explicit Finite-Volume Time-Domain (FV-TD) scheme for unstructured meshes. This will be performed within the GEMS project which is a collaborative code development project between Swedish industry and academia involving PSCI, Ericsson Saab Avionics, Ericsson Microwave System, Saab Ericsson Space, the Swedish Institute of Applied Mathematics and the Swedish Defense Research Establishment.

5.2 Parallel implementation

The parallel implementation of `pscyee` uses the Message Passing Interface (MPI) standard. This was chosen to guarantee portability of the code.

Figure 5.1 illustrates our parallelization strategy in 1D with six cells on two nodes. When node two calculates the Ex_4 -value it needs to know the value of $H_{z_{3.5}}$ which is stored in node one. Similarly, when node one calculates the $H_{z_{3.5}}$ -value it needs to know the value of Ex_4 which is stored in node two. This means that during each time step two messages have to be sent, one in each direction. In 3D two of the six electromagnetic variables must be sent across processor interfaces. We have implemented this strategy with several syntactically different MPI implementations. They are described in detail in [4].




Figure 5.1. Illustration of our parallelization strategy. A 1D Yee grid with six cells ($n=6$) and $\Delta y = 0.5$ distributed on two nodes.

On a parallel computer with reasonably fast communication, e.g. an IBM SP, it is possible to achieve perfect scalability. With perfect scalability we mean that when the problem size is increased linearly with the number of processors the execution time is constant. On the other hand, it is usually not possible to achieve ideal speed-up. An exception occurs when cache effects makes smaller problem sizes faster than larger ones. With speed-up we mean that the problem size is kept fixed independently of the number of processors. Figure 5.2 displays speed-up results for

an IBM SP. A point source was used for excitation and the first-order Mur scheme was used as ABC. The nonblocking communication of MPI (MPI_ISEND) was used.




Figure 5.2. Speed-up of `pscyeon` on an IBMSP with 160MHz RS/6000 processors. The problem size is $252 \times 252 \times 127$. The solid line represents ideal speed-up.

Portability has been demonstrated by executing the code on several different parallel computers, including an IBM SP, a Fujitsu VX, a CRAY J932 and a cluster of Digital Alpha servers. The performance was satisfactory on all these computers.

A more detailed descriptions of our parallelization strategy can be found in [3, 4]. A nice review of parallelization of FD-TD is given by Professor Gedney in Chapter 5 of [52].

5.3 One billion cells

For homogeneous materials we need only six floating point values per cell, the three electric and the three magnetic field components. Using 32-bit precision (four bytes) means that we need $4 * 6 * 10^9 = 22.4$ Gbyte memory for one billion cells. Clearly, only a parallel computer can supply this. We used 125 nodes with 160MHz RS/6000 processors of the IBM SP at PDC, KTH. These nodes have 256 Mbyte memory each making a total of 31.25 Gbyte. Actually, a few of them have more than 256 Mbyte, but we do not want to use it since that would destroy the load balancing.

It is not possible to use all the physical memory on a node since some of the memory is used by the operating system. Using too much memory will result in swapping which must be avoided since it has a drastic effect on the performance. Tests indicated that it was safe to use up to 200 Mbyte on each node and that one usually could use up to 220 Mbyte.

The object chosen for this calculation was a SAAB 2000 aircraft. A Cartesian description of this aircraft was created from a CAD description using CADfix. This was performed by Ericsson Saab Avionics (ESB). The file delivered from ESB only contained one half of the airplane so we had to create the entire airplane by mirroring. Hence, we got an absolutely symmetric airplane. The resolution was 2.5 cm




Figure 5.3. The surface currents after 1500 time steps on the SAAB 2000 aircraft. Also the magnitude of the \mathbf{H} -field is shown on a cutting plane across the wings perpendicular to the fuselage. (A color version is displayed on Page 172.)

in all three dimensions which was a factor two per space dimension smaller than the resolution previously used at ESB. We used a problem size of $1260 \times 1260 \times 635$ equaling 1 008 126 000 cells in total. The computational domain was split in $5 \times 5 \times 5$ blocks each with a size of $252 \times 252 \times 127$ cells.

The total number of \mathbf{E} -fields on the surface of the SAAB 2000 was almost two millions and the number of surface quads was almost one million. The memory needed to store this information varied from node to node and was at most nine Mbytes.

The input file containing the PEC information for one half of the SAAB 2000 contained almost one million lines. Since all nodes read this it was very inefficient to use the standard file system, `afs`, which only gave a CPU activity of 2–3% while reading the file. Instead we used the parallel file system, `pfs`, which gave a CPU activity of about 20%. A more efficient strategy would probably be to let one node read the data, analyze it and then distribute it to the other nodes.

The performance of the core of the code is almost 25 Gflop/s. This figure excludes output. When the surface currents and two cutting planes were written every 20th time step this performance dropped to about 20 Gflop/s. In this case, the time to complete one time step was about 1.8 seconds. A total of 2500 time steps were taken and the total execution time, including initializations, was 86 minutes.

The first-order Mur ABC was used and excitation was performed with a point source in front of the airplane. The `MPI_SENDRECV` facility was used for the communication.

5.4 Visualization of large FD-TD data




Figure 5.4. The interior of the SAAB 2000 aircraft. Surface currents are shown at the same time as in Figure 5.3. (A color version is displayed on Page 172.)

Running large FD-TD simulations do not only require powerful computers. It also put high demands on the post processors. For small and medium size problems you can save the whole electromagnetic field at each time step (or every n :th time step) if you have sufficiently large discs. After the simulation you can go through the data and visualize the features you are looking for and also find unexpected properties of your solution. This is not the case when you are solving large problems. Not only the disc space is limited but the I/O bandwidth is also an effective bottleneck. You have to decide a priori what field values you want to post process and save only them. Furthermore, you need a high end graphical system to visualize the multitude of polygons that constitute the objects in your simulation.

A new technique has emerged during the last few years which is believed to make the understanding of scientific computing results easier. The concept is usually called “CAVE”, CAVE Automated Virtual Environment, and consist of back-projection of images onto semitransparent surfaces. If several surfaces are put together you get a room where you are surrounded by the images. By adding a tracking system where your head position is tracked, stereo images can be produced and highly realistic 3D environments are perceived. With a tracking system for a hand held device you can also interact with this virtual reality.

This technique has obvious benefits: it is easy for several people to be in the room simultaneously and therefore see images together. The users can thus interact with the virtual reality together and focus on interesting areas. Also, because users

see their own hands and feet, for example, as part of the virtual world, they get a heightened sense of being inside that world.

Various CAVE-like environments exist all over the world today. Most of these have up to four projection surfaces; images are usually projected on three walls and the floor. Adding projection on the ceiling gives a fuller sense of being enclosed in the virtual world.

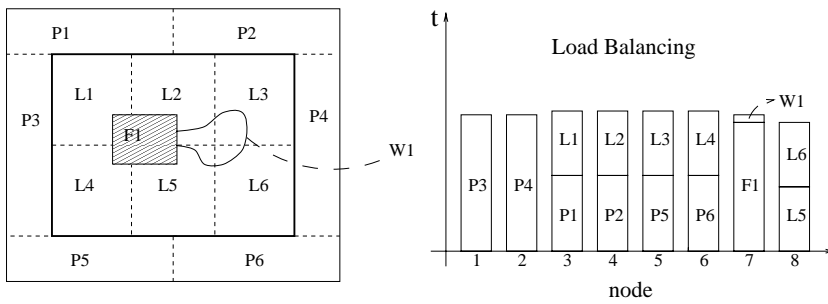
Projection on all six surfaces of a room allows users to turn around and look in all directions. Thus, their perception and experience are not limited, which is necessary for full immersion. Such a six-surface-system has recently been inaugurated at PDC, KTH and several projects on visualizing CEM solutions have been started where the users will be able to navigate, for example inside an aircraft while lightning strikes. In this case one will directly see the field penetrating the openings of the fuselage and detect “hotspots” to avoid in the context of EMC.

However, the CAVE technology does not ease the urge of effective handling of the output from FD-TD solvers. Even though the computers serving CAVE environments often are high end graphical systems you still have to limit the data saved for post processing. In Figure 5.3 the surface currents are displayed on each $2.5 \times 2.5 \text{ cm}^2$ square constituting the surface of the FD-TD object. Approximately one million quads are put to the visual system and clearly, most of them are not visible in the picture. Furthermore, perhaps one could utilize the concept “level of details” where smaller parts in the background are combined to fewer objects and thus lower the number of polygons to be rendered. For volumetric data semitransparent 3D texture mapping can be utilized. This volume rendering technique can be used to show the field inside the aircraft.

5.5 Conclusions

We have demonstrated that by using a large parallel computer it is possible to do an FD-TD calculation with more than one billion cells. We showed this by using 125 nodes of an IBM SP.

The parallel implementations of PML and the subcell model for arbitrarily oriented thin wires (Paper 5 on Page 10) are in progress. When they are completed we plan to repeat the one billion cells calculation using these features. Excitation will be performed using the thin-wire model. The parallelization approach is illustrated in Figure 5.5, where models for frequency dispersive materials are also included.



P = PML, L = Leap frog, W = Wire, F = Freq. disp. mat.

Figure 5.5. The load balancing for a full problem, including PML, frequency dispersive materials and an arbitrarily oriented thin-wire subcell model.

5.6 Acknowledgments

This work was supported with computing resources by the Swedish Council for Planning and Coordination of Research (FRN) and Paralleldatorcentrum (PDC), Royal Institute of Technology.

The grid for the SAAB 2000 was supplied by Ericsson Saab Avionics.

Chapter 6

Hybrid Methods in 2D

6.1 Introduction

In this section we present a detailed description and a thorough performance analysis of our hybrid methods in 2D. We consider the TM equations given in (3.8) with $\sigma = \sigma^* = 0$. We demonstrate that staircasing of a circular cylinder destroys the second-order accuracy of FD-TD and that second-order accuracy is recovered when using our hybrid methods. We also present numerical stability experiments.

This Chapter is based on Paper 2 in the list in Chapter 1. Results for discontinuous μ and ϵ have been added. There has also been some improvements in the notation, especially in the Finite-volume section. The description of the Finite-volume method is based on the the Licentiate thesis of Fredrik Edelvik [29].

6.2 Finite-difference method

The FD-TD method is described in Chapter 4. The TM grid is illustrated in Figure 4.1. Here we give a brief additional description. We recall that the fields are staggered both in time and space. This means that when the electric field E_z is known on a time-level, the magnetic field components H_x and H_y can be explicitly calculated on the next half time-level using only the previous H_x and H_y values and the latest E_z values, see Figure 6.1. We use the U-PML method described in Section 4.6 and the profile for σ is the one suggested in [38], i.e.

$$\sigma(\rho) = \frac{1}{3\pi d} \left(\frac{\rho}{d}\right)^4 \quad (6.1)$$

where d is the thickness of the PML layer, ρ is the distance to the interface between PML and the computational domain.




Figure 6.1. The timestepping mechanism for the FD-TD method: only H_y and E_z are considered here.

6.3 Finite-element method

6.3.1 FE-TD formulation

The FE-TD formulation is based on the second-order differential equation obtained by eliminating either the \mathbf{H} - or \mathbf{E} -field in the Maxwell equations. In the 2D TM case with $\sigma = \sigma^* = 0$ the \mathbf{E} -field is eliminated to yield

$$\mu \frac{\partial^2 \mathbf{H}}{\partial t^2} + \nabla \times \frac{1}{\epsilon} \nabla \times \mathbf{H} = \mathbf{0}. \quad (6.2)$$

Together with initial values and boundary conditions this defines the problem to solve. As boundary conditions (BC) we use:

$$\mathbf{n} \times (\nabla \times \mathbf{H}) = \mathbf{0} \quad \text{on } \Gamma_e \text{ (PEC)}, \quad (6.3)$$

$$\mathbf{n} \times \mathbf{H} = \mathbf{0} \quad \text{on } \Gamma_h \text{ (PMC)}, \quad (6.4)$$

$$\mathbf{n} \times \mathbf{H} = \mathbf{n} \times \mathbf{H}_{fddd}(t) \quad \text{on } \Gamma_t. \quad (6.5)$$

In the TM case the first type of BC (PEC) becomes a Neumann boundary condition. The second type models a PMC and becomes a Dirichlet boundary condition. The third type is also a Dirichlet boundary condition, but it is time dependent, and is used at the interface to the FD-TD domain.

For the hybridization we also integrate E_z in the transition layer between the FE-TD domain and FD-TD domain by using Ampère's law. These E_z components are required by the FD-TD scheme as boundary values, see Section 6.6. For visualization purposes E_z is actually integrated in the whole domain.

6.3.2 Spatial discretization

The weak form, or the Galerkin form, of our problem can be stated as: find $\mathbf{H} \in W$, where $W = H(\text{curl}, \Omega) = \{v : v \in L^2(\Omega), \nabla \times v \in L^2(\Omega)\}$, such that

$$\int_{\Omega} \left(\mu \frac{\partial^2 \mathbf{H}}{\partial t^2} \cdot \mathbf{w} + \frac{1}{\epsilon} \nabla \times \mathbf{H} \cdot \nabla \times \mathbf{w} \right) d\Omega = - \int_{\Gamma} \frac{1}{\epsilon} \mathbf{n} \times \nabla \times \mathbf{H} \cdot \mathbf{w} ds , \quad (6.6)$$

for all $\mathbf{w} \in W$. As trial and test functions we have chosen “edge” or “Whitney” elements [76]. These elements constitute the natural FEM analogue to the Yee scheme in that they essentially yield the same scheme on a Yee grid [69]. They give a “physical” approximation in the sense that only tangential continuity across element edges is enforced, and not normal continuity. This is in agreement with the interface condition given in (3.16). A comparison between edge elements and standard node elements have been presented in [68]. The edge elements have the advantage that the Gauss’ laws are better fulfilled by the approximations, and it is easy to implement essential boundary condition.

To define these linear edge elements which are second-order accurate, consider the standard linear basis functions Φ_i for nodal-based finite elements, constructed such that $\Phi_i = 1$ in node number i and $\Phi_i = 0$ in all other nodes. Take edge e to be the edge on a triangular element joining node i and node j . The basis function for edge e is

$$\varphi_e = \Phi_i \nabla \Phi_j - \Phi_j \nabla \Phi_i , \quad (6.7)$$

and is illustrated in Figure 6.2.




Figure 6.2. The basis function φ_e for edge e plotted over a triangle.

The basis function for edge e has the following properties:

- $\nabla \cdot \varphi_e = 0$.
- φ_e has constant tangential component ($= 1/\text{length}$) along edge e , which means that the tangential component is continuous around edge e . The normal component is discontinuous.
- φ_e has zero tangential component along the other edges.

The approximation $\mathbf{H}^h \in W^h$ of $\mathbf{H} \in W$ can then be written as $\mathbf{H}^h = \sum_e \alpha_e \varphi_e$, where α_e is an approximation of the \mathbf{H} component tangential to edge e , multiplied by the length of edge e . The edge element method applied to (6.6) then becomes: solve the system of ordinary differential equations

$$\mathbf{M} \frac{d^2 \alpha}{dt^2} + \mathbf{S} \alpha = \mathbf{f}^A, \quad (6.8)$$

where the vector α contains the unknown α_e :s. The matrix elements coupling edge k and edge l is:

$$(\mathbf{M})_{kl} = \int_{\Omega} \mu \varphi_k \cdot \varphi_l d\Omega, \quad (6.9)$$

$$(\mathbf{S})_{kl} = \int_{\Omega} \frac{1}{\epsilon} (\nabla \times \varphi_k) \cdot (\nabla \times \varphi_l) d\Omega, \quad (6.10)$$

and the right-hand side \mathbf{f}^A contains the contributions from the boundary conditions, i.e. we have $\mathbf{f}^A = \mathbf{f}^D + \mathbf{f}^N$. The general form of the contributions from a time dependent Dirichlet boundary Γ_t is

$$(\mathbf{f})_k^D = - \sum_{l \in \Gamma_t} \left((\mathbf{M})_{kl} \frac{\partial^2 \alpha_l}{\partial t^2} + (\mathbf{S})_{kl} \alpha_l \right). \quad (6.11)$$

The PMC BC is just a simple special case of this, where these terms become zero. In the hybrid code the “boundary values” from the FD-TD domain will enter the FE-TD domain in this way. At a Neumann boundary Γ_e , $\mathbf{n} \times \nabla \times \mathbf{H}$ is to be specified, see (6.3). If this expression is a known function on the boundary then

$$(\mathbf{f})_k^N = - \int_{\Gamma_e} \frac{1}{\epsilon} \mathbf{n} \times \nabla \times \mathbf{H} \cdot \varphi_k ds, \quad (6.12)$$

has to be computed.

When formulating the local mass and stiffness matrices it is convenient to introduce the following notation. Define

$$\mathbf{l}_i = (x_{i-1} - x_{i+1}, y_{i-1} - y_{i+1})^T, \quad (6.13)$$

as the vector along edge i , with direction counterclockwise (see Figure 6.3). The index i in (6.13) assumes the values 1, 2 and 3 cyclically, so that if $i = 3$ then $i + 1 = 1$. Furthermore, let the the vector \mathbf{d} be given by

$$\mathbf{d} = (d_1, d_2, d_3)^T, \quad (6.14)$$

where $d_i = 1$ if the local direction of edge vector \mathbf{l}_i is the same as the globally defined direction, otherwise $d_i = -1$. Expressed in this notation the local three by three geometrical stiffness matrix \mathbf{s}^g becomes

$$\mathbf{s}^g = \mathbf{d} \mathbf{d}^T / A, \quad (6.15)$$




Figure 6.3. Triangular element with nodes and vectors along edges.

where A is the area of the triangular element, and the geometrical mass matrix \mathbf{m}^g has the elements

$$\begin{aligned} m_{ii}^g &= (\mathbf{l}_{i+1}^T \mathbf{l}_{i+1} + \mathbf{l}_{i-1}^T \mathbf{l}_{i-1} - \mathbf{l}_{i+1}^T \mathbf{l}_{i-1}) / (24A), \\ m_{ij}^g &= -(\mathbf{l}_k^T \mathbf{l}_k + \mathbf{l}_i^T \mathbf{l}_j) / (24A \cdot d_i d_j), \quad i \neq j, \quad j \neq k, \quad k \neq i, \end{aligned} \quad (6.16)$$

where the indexes i, j, k are cyclically defined as explained above. The material properties μ and ϵ are assumed to be constant within each triangle. The matrices \mathbf{m} and \mathbf{s} are thus given by scalars times geometric mass/stiffness matrices, i.e.

$$\mathbf{m} = \mu \cdot \mathbf{m}^g, \quad (6.17)$$

$$\mathbf{s} = \frac{1}{\epsilon} \cdot \mathbf{s}^g. \quad (6.18)$$

6.3.3 Time discretization

A major drawback of the edge element method is that it is difficult to obtain an explicit scheme on a general unstructured grid. For the first-order formulation of the Maxwell equations, it has been shown that masslumping can be justified only if all the triangles are acute [70]. However, the FE solver is meant to be used to resolve fine geometrical features and in that case an explicit time integrator cannot be used anyway due to the small time step required. Using (6.11) the system (6.8) may be written as

$$\begin{cases} (\mathbf{M} \ \mathbf{M}_D) \begin{pmatrix} \frac{d^2 \alpha}{dt^2} \\ \frac{d^2 \alpha_D}{dt^2} \end{pmatrix} + (\mathbf{S} \ \mathbf{S}_D) \begin{pmatrix} \alpha \\ \alpha_D \end{pmatrix} = \mathbf{f}, \\ \alpha(0) = \alpha_0, \\ \frac{d\alpha(0)}{dt} = \mathbf{v}_0, \end{cases} \quad (6.19)$$

where \mathbf{M} and \mathbf{M}_D are mass matrices, \mathbf{S} and \mathbf{S}_D are stiffness matrices from the discretization of the double-curl operator and α is the magnetic field vector, which we will solve for. The vector α_D contains the given magnetic field at the Dirichlet boundary and the matrices with index D describe the influence from this boundary.

Note that the right-hand side now is $\mathbf{f} = \mathbf{f}^N$. To get a first-order system we introduce \mathbf{v} as the time derivative of α and rewrite the system as

$$\begin{cases} \begin{pmatrix} \frac{d\alpha}{dt} \\ \frac{d\alpha_D}{dt} \end{pmatrix} = \begin{pmatrix} \mathbf{v} \\ \mathbf{v}_D \end{pmatrix}, \\ (M \ M_D) \begin{pmatrix} \frac{d\mathbf{v}}{dt} \\ \frac{d\mathbf{v}_D}{dt} \end{pmatrix} + (S \ S_D) \begin{pmatrix} \alpha \\ \alpha_D \end{pmatrix} = \mathbf{f}. \end{cases} \quad (6.20)$$

For the time-discretization of this system we have implemented two methods, the two stage backward difference formula (BDF-2) and the θ -method. The θ -method yields

$$\frac{\alpha^{n+1} - \alpha^n}{\Delta t} = \theta \mathbf{v}^{n+1} + (1 - \theta) \mathbf{v}^n, \quad (6.21)$$

$$M \frac{\mathbf{v}^{n+1} - \mathbf{v}^n}{\Delta t} + S(\theta \alpha^{n+1} + (1 - \theta) \alpha^n) =$$

$$\theta \mathbf{f}^{n+1} + (1 - \theta) \mathbf{f}^n - M_D \frac{\mathbf{v}_D^{n+1} - \mathbf{v}_D^n}{\Delta t} - S_D(\theta \alpha_D^{n+1} + (1 - \theta) \alpha_D^n), \quad (6.22)$$

where

$$\mathbf{v}_D^{n+1} = \left(-(1 - \theta) \mathbf{v}_D^n + \frac{\alpha_D^{n+1} - \alpha_D^n}{\Delta t} \right) / \theta, \quad (6.23)$$

and α^n corresponds to $\alpha(t_n)$, $\Delta t = t_{n+1} - t_n$ and $0 < \theta \leq 1$. Note that we use the θ -method not only for the time-integration but also for approximating the time-derivative \mathbf{v}_D^{n+1} on the right-hand side of (6.22). The solution process in each time step begins by applying (6.23) to compute \mathbf{v}_D^{n+1} given the boundary condition α_D^{n+1} and then solve the system (6.22), which is rewritten as

$$\begin{bmatrix} I/\Delta t & -\theta I \\ \theta S & M/\Delta t \end{bmatrix} \begin{bmatrix} \delta\alpha \\ \delta\mathbf{v} \end{bmatrix} = - \begin{bmatrix} \mathbf{res}_\alpha \\ \mathbf{res}_v \end{bmatrix}, \quad (6.24)$$

where

$$\begin{aligned} \delta\alpha &= \alpha^{n+1} - \alpha^n, \\ \delta\mathbf{v} &= \mathbf{v}^{n+1} - \mathbf{v}^n, \end{aligned} \quad (6.25)$$

and the residuals on the right hand side are

$$\begin{aligned} \mathbf{res}_\alpha &= -\mathbf{v}^n, \\ \mathbf{res}_v &= S\alpha^n - \theta \mathbf{f}^{n+1} - (1 - \theta) \mathbf{f}^n \\ &+ M_D \frac{\mathbf{v}_D^{n+1} - \mathbf{v}_D^n}{\Delta t} + S_D(\theta \alpha_D^{n+1} + (1 - \theta) \alpha_D^n), \end{aligned} \quad (6.26)$$

The solution of (6.24) utilizes the following block LU factorization of the Jacobian of the system

$$\begin{bmatrix} \mathbf{I}/\Delta t & -\theta\mathbf{I} \\ \theta\mathbf{S} & \mathbf{M}/\Delta t \end{bmatrix} = \begin{bmatrix} \mathbf{I}/\Delta t & \mathbf{0} \\ \theta\mathbf{S} & \mathbf{J} \end{bmatrix} \begin{bmatrix} \mathbf{I} & -\Delta t\theta\mathbf{I} \\ \mathbf{0} & \mathbf{I} \end{bmatrix}, \quad (6.27)$$

where the Schur complement matrix \mathbf{J} is given by

$$\mathbf{J} = \mathbf{M}/\Delta t + \Delta t\theta^2\mathbf{S}. \quad (6.28)$$

The solution can then be divided into the three steps

$$\begin{aligned} \delta\alpha^p &= -\Delta t\mathbf{res}_\alpha, \\ \mathbf{J}\delta\mathbf{v} &= -\mathbf{res}_v - \Delta t\theta\mathbf{S}\delta\alpha^p, \\ \delta\alpha &= \delta\alpha^p + \Delta t\theta\delta\mathbf{v}. \end{aligned} \quad (6.29)$$

The first two steps are the forward substitution in the block LU factorization while the third step is the back substitution.

This scheme is unconditionally stable for $1/2 \leq \theta \leq 1$. With $\theta = 1/2$ we obtain the only second-order θ -method, the midpoint or Crank-Nicholson method. Note that this method is equivalent to the method used by Hwang and Wu [51] on the second-order system. As an alternative we have implemented the BDF-2 method, which also is second-order accurate and unconditionally stable. Applied to

$$\frac{dx}{dt} = \mathbf{f}(x), \quad (6.30)$$

the BDF-2 method, with constant time step, becomes

$$\frac{3x^{n+1} - 4x^n + x^{n-1}}{2\Delta t} = \mathbf{f}(x^{n+1}). \quad (6.31)$$

The BDF-2 method has an error constant which is four times larger than the error constant of the Crank-Nicholson method, but it has shown better stability properties than the Crank-Nicholson method when used in the hybrid code. We note that the time-integration requires initial values for the “velocities”. This is similar to the requirement of initial data in the Yee scheme, and we simply set an initial state where all fields and their time derivatives are zero. Also note that BDF-2 requires values at two previous time steps, which for the initial values is equivalent to knowing also the initial “accelerations”. This would normally require some special treatment to start the integration, but since we assume zero initial values for the “accelerations” this poses no problem.

6.3.4 Workload and memory requirements

The workload per time step is of the order $4bN$ where b is the bandwidth of the system (6.24) after renumbering and N is the number of unknowns which is equal to the number of edges in the unstructured grid. Neglecting boundaries, the number of edges equal 1.5 times the number of triangles. The factorization takes $2b^2N$ arithmetical operations, but it is only performed once. The bandwidth is of course problem dependent. The memory requirements are bN floating point values. We use the reverse Cuthill-McKee [41] ordering algorithm to reduce the bandwidth of the system (6.24).

The use of an iterative method would decrease the computational workload and the memory requirements. Iterative solvers will be further discussed in Chapter 7.

6.4 Finite-volume method

6.4.1 FV formulation

The Finite-volume (FV) solver is based on the following integral formulation of Ampère's and Faraday's laws:

$$\frac{\partial}{\partial t} \int_A \mu \mathbf{H} dA = - \oint_{\Gamma} \mathbf{n} \times \mathbf{E} dl, \quad (6.32)$$

$$\frac{\partial}{\partial t} \int_A \epsilon \mathbf{E} dA = \oint_{\Gamma} \mathbf{n} \times \mathbf{H} dl, \quad (6.33)$$

where A is an arbitrary area, Γ is the path that encloses A and \mathbf{n} is the unit normal.

6.4.2 Spatial discretization

The integral formulations in (6.32) and (6.33) are discretized on a staggered grid by introducing a dual grid to the primary triangular grid. The magnetic components are situated at the nodes of the primary grid and the electric components, in the 2D TM case only E_z , are situated at the nodes of the dual grid. The dual grid is created at the preprocessing stage by defining dual nodes at the barycenters of the primary cells; see Section 6.4.5 for a detailed description.

In the 2D TM case integrating the magnetic and electric fields over each dual and primary cell gives the following integral form:

$$\frac{\partial}{\partial t} \int_{A_j^d} \tilde{\mu}_j^d \mathbf{H} dA = - \sum_k \int_{\Gamma_{j,k}^d} \mathbf{n}_{j,k}^d \times (E_z \hat{\mathbf{z}}) dl, \quad (6.34)$$

$$\frac{\partial}{\partial t} \int_{A_i^p} \epsilon_i^p E_z dA = \sum_m \int_{\Gamma_{i,m}^p} \mathbf{n}_{i,m}^p \times \mathbf{H} dl, \quad (6.35)$$

where A_j^d is the area of dual cell j , Γ_j^d is the path that encloses A_j^d and $\mathbf{n}_{j,k}^d$ are the unit edge normals for the dual edges k in dual cell j . The variables belonging to the primary cell i are defined similarly.

All materials are defined relative to primary grid cells. For dielectric materials that is a natural definition. However, the magnetic permeability, μ , is associated with the dual grid cells. Therefore, μ requires averaging. The average permeability of dual cell j is computed as

$$\tilde{\mu}_j^d = \frac{1}{A_j^d} \sum_q \mu_q^p A_{j,q}^d, \quad (6.36)$$

where $A_{j,q}^d$ is the part of the area A_j^d that is inside primary cell q . Performing the material averaging in this manner preserves the second-order accuracy of the solver for inhomogeneous materials.

The area integrals in (6.34) and (6.35) are evaluated by taking the average values of the fields multiplied by the areas of the respective cells. Simplifying the two integrands in the line integrals implies

$$\tilde{\mu}_j^d \frac{\partial}{\partial t} \mathbf{H}_j = \frac{1}{A_j^d} \sum_k \int_{\Gamma_{j,k}^d} \tilde{E}_z|_{j,k} \mathbf{t}_{j,k}^d dl, \quad (6.37)$$

$$\epsilon_i^p \frac{\partial}{\partial t} E_z|_i = \frac{1}{A_i^p} \sum_m \int_{\Gamma_{i,m}^p} \tilde{\mathbf{H}} \cdot \mathbf{t}_{i,m}^p dl, \quad (6.38)$$

where $\mathbf{t}_{j,k}^d$ are unit vectors for the dual edges k in dual cell j and $\mathbf{t}_{i,m}^p$ are unit vectors for the edges m in primary cell i . The line integral in (6.37) is evaluated by assuming that the electric field is piecewise linear along the dual edges. Hence, $\tilde{E}_z|_{j,k}$ is computed by taking the arithmetic mean value of the electric field at the two nodes defining the dual edge, $\mathbf{t}_{j,k}^d$. However, we cannot use the same approach when calculating the integral in (6.38) since that does not guarantee that the divergence is preserved on a local cell level. This has been found to be very critical if spurious modes in the numerical solution are to be suppressed. The divergence is preserved if we incorporate an ‘‘FD-TD’’-correction along the edges in the primary grid (see Section 6.4.4 for a proof). Following Riley et al. [84], the magnetic field projected along the primary edge $\mathbf{t}_{i,m}^p$ is evaluated as

$$\begin{aligned} \tilde{\mathbf{H}} \cdot \mathbf{t}_{i,m}^p &= \mathbf{H} \cdot \mathbf{n}_{j,k}^d (\mathbf{n}_{j,k}^d \cdot \mathbf{t}_{i,m}^p) \\ &+ \frac{1}{2} [(\mathbf{H}_j + \mathbf{H}_r) - ((\mathbf{H}_j + \mathbf{H}_r) \cdot \mathbf{n}_{j,k}^d) \mathbf{n}_{j,k}^d] \cdot \mathbf{t}_{i,m}^p, \end{aligned} \quad (6.39)$$

where \mathbf{H}_j and \mathbf{H}_r are the magnetic field at the two nodes defining the primary edge and $\mathbf{H} \cdot \mathbf{n}_{j,k}^d$ is the FD-TD component in the direction $\mathbf{n}_{j,k}^d$ orthogonal to the

dual edge $\mathbf{t}_{j,k}^d$, which crosses the primary edge $\mathbf{t}_{i,m}^p$, see Figure 6.4. The FD-TD component is updated according to

$$\frac{\partial \mathbf{H} \cdot \mathbf{n}_{j,k}^d}{\partial t} = \frac{E_z|_q - E_z|i}{\Delta_{j,k}^d \tilde{\mu}_{j,k}^d}, \quad (6.40)$$

where $\Delta_{j,k}^d$ is the length and $\tilde{\mu}_{j,k}^d$ is the average permeability of the dual edge $\mathbf{t}_{j,k}^d$. The average permeability, $\tilde{\mu}_{j,k}^d$, is approximated as

$$\tilde{\mu}_{j,k}^d = \frac{\mu_q^p \Delta_q + \mu_i^p (\Delta_{j,k}^d - \Delta_q)}{\Delta_{j,k}^d}, \quad (6.41)$$

where i and q are the two primary cells sharing dual edge $\mathbf{t}_{j,k}^d$, Δ_q is the part of the length of the dual edge that is inside primary cell q , see Figure 6.4.




Figure 6.4. The magnetic field along $\mathbf{t}_{i,m}^p$ is corrected using the FD-TD value in direction $\mathbf{n}_{j,k}^d$.

Taking a closer look at (6.39) we see that if the primary and dual edges are orthogonal, the vectors $\mathbf{n}_{j,k}^d$ and $\mathbf{t}_{i,m}^p$ are parallel and the second part of (6.39) vanishes. Hence, the name “FD-TD”-correction is somewhat misleading since that term is actually the important one. The magnetic node values are only used to give a better approximation of the edge-projected field in the case when $\mathbf{n}_{j,k}^d$ and $\mathbf{t}_{i,m}^p$ do not align.

The boundary condition for a perfect magnetic conductor (PMC) gives us that the tangential component of the magnetic field at the object is zero. A complication occurs whenever the computation of $\tilde{\mathbf{H}} \cdot \mathbf{t}_{i,m}^p$ does not reduce to the FD-TD term, $\mathbf{H} \cdot \mathbf{n}_{j,k}^d$, where $\mathbf{t}_{i,m}^p$ denotes a primary edge with one node on the boundary. When that is not the case the following alternative is used (see Figure 6.5(a)):

$$\begin{aligned} \tilde{\mathbf{H}} \cdot \mathbf{t}_{i,m}^p &= \mathbf{H} \cdot \mathbf{n}_{j,k}^d (\mathbf{n}_{j,k}^d \cdot \mathbf{t}_{i,m}^p) \\ &+ \frac{1}{2} [\mathbf{H}_j + (\mathbf{H}_j \cdot \tilde{\mathbf{n}}_r) \tilde{\mathbf{n}}_r - (\mathbf{H}_j + (\mathbf{H}_j \cdot \tilde{\mathbf{n}}_r) \tilde{\mathbf{n}}_r) \cdot \mathbf{n}_{j,k}^d] \cdot \mathbf{t}_{i,m}^p, \end{aligned} \quad (6.42)$$

where $\tilde{\mathbf{n}}_r$ denotes an average normal at the boundary node. This normal is defined by taking the average direction of the two boundary edges including the boundary node and then taking the cross product with that direction and \hat{z} .




Figure 6.5. Primary and dual cells at PMC and PEC boundaries. The boundaries are illustrated by thick lines. The PEC boundary condition is enforced using the method of images.

For a perfect electric conductor (PEC) the tangential electric components, in the TM case only E_z , should equal zero at the boundary. However, the E_z components are not situated on the boundary. The boundary condition is enforced by setting the electric field inside the conductor equal to the value of the electric field directly outside the conductor with a change of sign, see Figure 6.5(b). These uniquely determined virtual image nodes are identified when the dual grid is constructed in the preprocessing phase, see Section 6.4.5.

6.4.3 Time discretization

We use a third-order staggered Adams–Bashforth scheme (ABS3) [40],

$$\mathbf{H}_j^{n+\frac{1}{2}} = \mathbf{H}_j^{n-\frac{1}{2}} \quad (6.43)$$

$$+ \frac{\Delta t}{\tilde{\mu}_j^d A_j^d} \sum_k \left(\frac{25}{24} \tilde{E}_z|_{j,k}^n - \frac{1}{12} \tilde{E}_z|_{j,k}^{n-1} + \frac{1}{24} \tilde{E}_z|_{j,k}^{n-2} \right) \mathbf{t}_{j,k}^d \Delta_{j,k}^d, \quad (6.44)$$

$$E_z|_i^{n+1} = E_z|_i^n + \frac{\Delta t}{\epsilon_i^p A_i^p} \sum_m \left(\frac{25}{24} \tilde{\mathbf{H}} \cdot \mathbf{t}_{i,m}^p|^{n+\frac{1}{2}} - \right. \\ \left. \frac{1}{12} \tilde{\mathbf{H}} \cdot \mathbf{t}_{i,m}^p|^{n-\frac{1}{2}} + \frac{1}{24} \tilde{\mathbf{H}} \cdot \mathbf{t}_{i,m}^p|^{n-\frac{3}{2}} \right) \Delta_{i,m}^p,$$

$$\mathbf{H} \cdot \mathbf{n}_{j,k}^d|^{n+\frac{1}{2}} = \mathbf{H} \cdot \mathbf{n}_{j,k}^d|^{n-\frac{1}{2}} + \Delta t \mathcal{F} \left(\frac{25}{24} E_z^n - \frac{1}{12} E_z^{n-1} + \frac{1}{24} E_z^{n-2} \right), \quad (6.45)$$

where \mathcal{F} is an operator taking care of the update of $\mathbf{H} \cdot \mathbf{n}_{j,k}^d$ according to (6.40). Since ABS3 is a staggered time integrator the time-coupling with FD-TD is straightforward, see Section 6.6. Furthermore, its stability properties are superior compared to the traditionally used leap-frog scheme, see Section 6.4.7.

6.4.4 Preservation of divergence

The finite-volume solver is based on an integral formulation of Ampère’s and Faraday’s laws. However, the Maxwell equations also include the Gauss’ laws that have to be satisfied in order to ensure a physical solution. For lossless materials, the divergence of the magnetic flux density vector, \mathbf{B} , as well as that of the electric flux density vector, \mathbf{D} , should equal zero.

Proposition 6.1. *For the 2D TM_z Maxwell equations, $\nabla \cdot \mathbf{D} = 0$ is automatically satisfied.*

Proof The proof is trivial because \mathbf{E} only has a z component in the TM case and the last term is zero because we do not have any variation in the z -direction in 2D,

$$\nabla \cdot \mathbf{D} = \nabla \cdot \epsilon \mathbf{E} = \frac{\partial \epsilon E_x}{\partial x} + \frac{\partial \epsilon E_y}{\partial y} + \frac{\partial \epsilon E_z}{\partial z} = 0.$$

□

Proposition 6.2. *The FV solver preserves $\nabla \cdot \mathbf{B}$ on a local cell level to machine precision.*

Proof For dual cell j we have

$$\frac{\partial}{\partial t} \int_{A_j^d} \nabla \cdot \mathbf{B} \, dA = \frac{\partial}{\partial t} \int_{A_j^d} \nabla \cdot \mu \mathbf{H} \, dA = \frac{\partial}{\partial t} \oint_{\Gamma_j^d} \mu \mathbf{H} \cdot \mathbf{n}_j^d \, dl,$$

where Gauss' theorem is used to get the last equality. Splitting the integral path into k segments gives us

$$\frac{\partial}{\partial t} \sum_k \int_{\Gamma_{j,k}^d} \mu \mathbf{H} \cdot \mathbf{n}_{j,k}^d \, dl.$$

So far we are still using continuous variables. If we think of the segments as being the dual edges and also assume that $\mathbf{H} \cdot \mathbf{n}_{j,k}^d$ and μ are constant along each edge we get

$$\frac{\partial}{\partial t} \sum_k \tilde{\mu}_{j,k}^d \mathbf{H} \cdot \mathbf{n}_{j,k}^d \int_{\Gamma_{j,k}^d} dl = \frac{\partial}{\partial t} \sum_k \tilde{\mu}_{j,k}^d \mathbf{H} \cdot \mathbf{n}_{j,k}^d \Delta_{j,k}^d,$$

where $\Delta_{j,k}^d$ is the edge length. The permeability is time independent and we can swap the sum and time derivative operator to obtain

$$\sum_k \tilde{\mu}_{j,k}^d \Delta_{j,k}^d \frac{\partial \mathbf{H} \cdot \mathbf{n}_{j,k}^d}{\partial t}.$$

Finally we can use (6.40) which gives us

$$\sum_k (E_z|_q - E_z|_i) = 0,$$

where the last expression equals zero because in the sum over k , each electric node value occurs twice with opposite signs.

Thus, the time derivative of the divergence is equal to zero and hence the divergence is preserved to machine precision in each cell. \square

Note that if $\mathbf{H} \cdot \mathbf{n}_{j,k}^d$ had been computed as a projection of the arithmetic mean value of the two magnetic node values (see Figure 6.4), instead of being updated using (6.40), the divergence would in general have differed from zero.

We would also like to point out that the proof is independent of how $\tilde{\mu}_{j,k}^d$ is calculated. However, if we use (6.41) in (6.40) and (6.36) in (6.37) we obtain a second-order accurate FV solver.

6.4.5 Creating the dual grid

Perhaps the most demanding task in implementing the FV solver is the construction of the dual grid from the primary grid. This is done in the preprocessing phase by joining barycenters of the neighboring cells. Hence, the dual grid is uniquely determined by the primary grid.

Starting with the inner edges, i.e. edges not situated on a boundary, each inner edge in the primary grid is crossed by a dual edge, see Figure 6.6. These dual edges belong to the dual elements surrounding the start- and stop-nodes of the respective primary edges.

At a PEC boundary, as already mentioned, the electric field nodes are mirrored in the boundary. The dual edges crossing the boundary edges are assigned to the dual elements surrounding the start- and stop-nodes of the boundary edges. The dual edges lying inside the PEC object are constructed by joining the image nodes together and they are assigned to the dual elements that hold both the dual nodes of these edges.




Figure 6.6. Construction of the dual grid from the primary grid.

Note that if we have a very thin PEC object it is possible that two dual elements lying on opposite sides of the boundary will overlap each other geometrically. This does not pose any problems because the image nodes are only virtual and the electric fields in these nodes are not updated, but assigned values from the field in their corresponding node.

The hybrid boundary is slightly more complicated. The nodes lying on a hybrid interface should be updated as usual, i.e. they have dual cells surrounding them, see Section 6.6. The dual edges that cross the edges lying on the interface boundary belong as usual to the dual elements surrounding the start- and stop-nodes of the boundary edges, see Figure 6.6. Generating the outermost dual edges in the interface layer is the trickiest part. Mainly because the interface is allowed to be staircased to minimize the number of unstructured cells. Hence, we cannot, as in

the PEC case, create the outermost dual edges by simply joining the dual nodes lying in the interface layer. We have to take care of the convex corners, which are always four more than the concave corners. We know that we have a convex corner when the midpoint of the dual edge joining two consecutive dual nodes, moving counterclockwise in the interface layer, is the same point as the corner node in the primary grid; see the dotted line in bottom right corner of Figure 6.6. If that is the case we construct a new dual node and remove the dual edge going through the corner. Instead we generate the two dual edges going to and from the new node.

When all dual nodes and dual edges are constructed, the edges in each dual element are sorted such that they are traversed counterclockwise for each dual element. The area of a dual element is calculated by dividing it into triangles and summing the contributions from each triangle.

6.4.6 Workload and memory requirements

To implement the solver efficiently we have chosen to compute as much as possible initially. Hence, the update of the fields is accomplished using matrices acting on the respective vectors. Due to the fact that the matrices will be sparse but without structure we store them in compress sparse row format. This format is a very memory efficient format for sparse matrices and gives fast access to the matrix elements. After the matrices have been created, all grid variables can be written to disk. Hence, this approach is much more efficient than recomputing the expressions needed to update the field variables at every time step or using indirect addressing in several levels.

To be able to obtain a theoretical estimate of the memory requirements and number of arithmetic operations of the solver we have to make a few assumptions. First of all we neglect boundary conditions. The following is assumed about the grid:

- There are n triangles.
- There are three edges in each triangle and each edge is shared by two triangles.
 \Rightarrow There are $1.5n$ edges.
- There are three nodes in each triangle and each node belongs to six triangles.
 \Rightarrow There are $0.5n$ nodes and six dual edges per dual cell.

The field variables that we need to store are E_z , $\mathbf{H} \cdot \mathbf{n}^d$, \mathbf{H} and $\tilde{\mathbf{H}} \cdot \mathbf{t}^p$. Using ABS3, E_z and $\tilde{\mathbf{H}} \cdot \mathbf{t}^p$ are stored at three time levels, whereas $\mathbf{H} \cdot \mathbf{n}^d$ and \mathbf{H} only are stored at the latest time level.

Variable	Floating point numbers
E_z	$3 \cdot n$
$\mathbf{H} \cdot \mathbf{n}^d$	$1.5n$
\mathbf{H}	$2 \cdot 0.5n$
$\tilde{\mathbf{H}} \cdot \mathbf{t}^p$	$3 \cdot 1.5n$
Σ	$10n$

Table 6.1. Memory requirements to store the field variables.

The memory requirements for the matrix operators used to update the field variables can be calculated from (6.37)–(6.40).

Variable	Floating point numbers	Integers
E_z	$n \cdot 3 = 3n$	$3n + n = 4n$
$\mathbf{H} \cdot \mathbf{n}^d$	$1.5n \cdot 2 = 3n$	$3n + 1.5n = 4.5n$
\mathbf{H}	$n \cdot 6 = 6n$	$6n + n = 7n$
$\tilde{\mathbf{H}} \cdot \mathbf{t}^p$	$1.5n \cdot 5 = 7.5n$	$7.5n + 1.5n = 9n$
Σ	$19.5n$	$24.5n$

Table 6.2. Memory requirements to store the matrices.

Hence, using the results in Tables 6.1 and 6.2 the total memory requirements for the FV solver is 29.5 floating point numbers and 24.5 integers per cell. Using 64-bit precision for the floating point numbers and 32-bit precision for the integers this means that approximately 334 bytes per cell are needed.

The number of arithmetic operations used by the solver are calculated from (6.37)–(6.40) and (6.43)–(6.45).

Variable	Arithmetic operations
E_z	$n \cdot 11 = 11n$
$\mathbf{H} \cdot \mathbf{n}^d$	$1.5n \cdot 4 = 6n$
\mathbf{H}	$n \cdot 12 = 12n$
$\tilde{\mathbf{H}} \cdot \mathbf{t}^p$	$1.5n \cdot 14 = 21n$
\tilde{E}_z	$n \cdot 5 = 5n$
Σ	$55n$

Table 6.3. Number of arithmetic operations per cell.

Hence, from Table 6.3 we conclude that the FV solver performs approximately 55 arithmetic operations per cell for each time step. This should be compared with FD-TD, which needs eleven arithmetic operations per cell and 36 bytes of memory. Hence, the FV solver is a factor 5 more expensive in terms of arithmetic operation per cell and a factor of nine in terms of memory per cell compared to FD-TD.

Note that in the above calculations we have not used the fact that when the primary and dual edges are orthogonal the second part of (6.39) vanishes and the \mathbf{H} values are not needed. A case where this will happen is the equilateral grid. An implementation of the solver should of course take advantage of this and the memory requirements and workload will go down considerably. On an equilateral grid, or any other grid where orthogonality occurs, the memory requirements are 16.5 floating point numbers and 10 integers per cell. The total memory requirement is in this case 172 bytes using the same precision. The number of arithmetic operations goes down to 31 per cell for such a grid. Thus, it is possible to reduce the memory requirements and the number of arithmetic operations by approximately a factor of two for the special case when the grids are mutually orthogonal.

6.4.7 Stability analysis

We will now investigate the stability properties of our FV scheme. The stability region for ABS3 for the scalar test equation $u' = \lambda u$ is given in Figure 6.7. The




Figure 6.7. Stability region for ABS3.

scheme is stable for λ between $\pm 12/7$ along the imaginary axis compared to the leap-frog scheme, which is stable between ± 2 . That implies that we have to use a shorter time step for ABS3. However, the main disadvantage with the leap-frog scheme is that it is only stable on the imaginary axis and becomes unstable as soon as we have eigenvalues with a nonzero real part, which we are likely to have on unstructured grids and when boundaries are taken into account.

By introducing operators \mathcal{A} and \mathcal{B} that take care of the spatial discretization we are able to write the semi-discrete problem on matrix form as

$$\frac{\partial}{\partial t} \begin{pmatrix} \hat{\mathbf{H}} \\ E_z \end{pmatrix} = \begin{pmatrix} \mathbf{0} & \mathcal{A} \\ \mathcal{B} & \mathbf{0} \end{pmatrix} \begin{pmatrix} \hat{\mathbf{H}} \\ E_z \end{pmatrix}, \quad (6.46)$$

where $\hat{\mathbf{H}} = (\mathbf{H} \quad \mathbf{H} \cdot \mathbf{n}^d)$.




Figure 6.8. Three different uniform triangular grids, the equilateral grid (left), the one-directional grid (middle) and the diamond grid (right). The respective dual grids are indicated by dashed lines.

A detailed Fourier analysis of the stability and dispersion properties on the three types of unstructured grids in Figure 6.8 is given in the Licentiate thesis of Fredrik Edelvik [29]. This analysis shows that the stability condition for leap-frog timestepping is given by

$$c\Delta t_{LF} \leq \frac{\Delta}{\sqrt{2}} \cdot \frac{\Delta_{min}}{\Delta}, \quad (6.47)$$

where the first part in the right-hand side is the stability condition on Cartesian grids and Δ_{min} equals the shortest edge length in the primary and dual grids. The stability condition for ABS3 timestepping is given by

$$c\Delta t_{ABS3} \leq \frac{6}{7} \frac{\Delta}{\sqrt{2}} \cdot \frac{\Delta_{min}}{\Delta}. \quad (6.48)$$

To analyze the eigenvalues for a general unstructured grid including its boundaries we can no longer use Fourier analysis. Instead, let

$$\mathbf{z}^n = \left(E_z|^n \quad \hat{\mathbf{H}}^{n-\frac{1}{2}} \quad E_z|^{n-1} \quad \hat{\mathbf{H}}^{n-\frac{3}{2}} \quad E_z|^{n-2} \right)^T, \quad (6.49)$$

and after some straightforward rearrangements we are able to write (6.43)–(6.45) on matrix form as $\mathbf{z}^{n+1} = \mathcal{P}(\mathcal{A}, \mathcal{B})\mathbf{z}^n$, where

$$\mathcal{P} = \begin{pmatrix} \mathcal{I} + \frac{625}{576} \Delta t^2 \mathcal{B}\mathcal{A} & \frac{23}{24} \Delta t \mathcal{B} & -\frac{25}{288} \Delta t^2 \mathcal{B}\mathcal{A} & \frac{1}{24} \Delta t \mathcal{B} & \frac{25}{576} \Delta t^2 \mathcal{B}\mathcal{A} \\ \frac{25}{24} \Delta t \mathcal{A} & \mathcal{I} & -\frac{1}{12} \Delta t \mathcal{A} & \mathbf{0} & \frac{1}{24} \Delta t \mathcal{A} \\ \mathcal{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathcal{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathcal{I} & \mathbf{0} & \mathbf{0} \end{pmatrix}, \quad (6.50)$$

where \mathcal{A} and \mathcal{B} are defined as in (6.46). Analyzing the eigenvalues of the companion matrix \mathcal{P} , for the grid shown in Figure 6.9, reveals that if we choose the time step

for ABS3 in the same manner as above we get the eigenvalue spectrum shown in Figure 6.9, where all eigenvalues are within the unit circle. If we use the leap-frog scheme with the same time step, the largest eigenvalue is of the order 1.0003. Hence the leap-frog scheme is unstable even for a time step well within the stability limit along the imaginary axis. The eigenvalues close to the origin in Figure 6.9 are the “parasitic” roots of the ABS3 multi-step scheme. However, since these roots are all close or equal to zero they are quickly damped away.




Figure 6.9. Eigenvalues of the companion matrix \mathcal{P} using ABS3 for a PMC cylinder scattering case. The primary grid around the cylinder is shown to the right.

The values of the elements of \mathcal{P} are taken directly from the code. Hence, we do not need analytical definitions of these elements to perform the eigenvalue analysis.

6.5 Grid requirements

The cell size of the structured Cartesian grid is chosen by balancing the computational cost of using smaller cells and the inability to capture the physics by using too large cells. When it comes to unstructured grids the cell size does not have to be constant throughout the whole grid. Smaller cells are used where fine geometrical details need to be resolved and larger cells are used as much as possible to save computer resources. If the cell size in the Cartesian grid is based on resolving characteristic wavelengths, then the unstructured cells must not be larger than the structured cells.

The local wave propagation speed for waves propagating through a discrete grid depends on the local grid size. If the grid size changes abruptly this will inevitably give spurious reflections. Even if the cell size is changed in a smooth way there will be some reflections although they can be made small. A few fundamental observations about unstructured grids can be made:

1. The rate of change in grid size must be controlled, see Figure 6.10a.
2. The cells should be as close to equilateral as possible in homogeneous domains, see Figure 6.10b.
3. At material interfaces the cell size could be changed abruptly by a factor of $\sqrt{\epsilon_1/\epsilon_2}$ (for a dielectric transition). This is due to the change in wave impedance for the physical problem, see Figure 6.10c.




Figure 6.10. a) Cell size changes must be reasonable. b) Equilateral cells are preferable. c) Cell sizes are only allowed to change abruptly at material interfaces.

Some comments about the three points above are necessary:

1. One would of course like to have a very smooth transition from large to small cells and vice versa, but this means that many unstructured cells are required. For performance reasons we would like to have as few unstructured cells as possible and hence a well balanced compromise must be found between these two contradictory desires. Moreover, a too rapid transition in the FV case leads to dual cells that are ill-balanced, i.e. a primary node will be surrounded by a dual cell whose center is far from the primary node.
2. The reasons for the desire for equilateral triangles is that skewed triangles augment the local numerical error. This is expressed as a stronger anisotropy of the local wave propagation speed.
3. Abrupt change of the cell size at material interfaces is possible due to the fact that the number of cells per wavelength is conserved in the normal direction. This must of course be balanced against the second observation above. However, in the tangential direction there is not much one can do since the nodes on the material interface are shared between the cells on both sides.

The unstructured grid should be as free as possible from global anisotropy, which means that the orientation of the individual cells should be as random as possible or counteract each others' directivity as much as possible. This desire comes from




Figure 6.11. a) The local wave propagation speed c depends on the propagation direction ($c_1 \neq c_2 \neq c_3$). b) Global directivity might amplify global anisotropy. c) Global anisotropy might be reduced by locally counteracting anisotropies.

the fact that the local numerical wave propagation speed is not constant for all angles of propagation. See Figure 6.11 for illustrations.

Coupling unstructured grids to structured Cartesian grids is described in Section 6.6. In this case the unstructured grids must have a transition layer of split rectangles which coincide with the Cartesian cells. An example of a transition layer is shown in the left part of the unstructured grid in Figure 6.12.

6.6 Hybridization

It is very important that an electromagnetic solver for realistic industrial applications can handle complex geometries without compromising the efficiency more than necessary. We believe that the best way to achieve this is to combine structured grids with unstructured grids thereby getting a hybrid method.

Our strategy of using hybrid techniques between unstructured grids and structured grids is based on the following observations:

- The FD-TD method is an extremely efficient second-order method for homogeneous materials, with respect to time and memory consumption.
- The Cartesian grid handles general geometries poorly, and fine details cannot be taken into account without special tricks.
- An unstructured grid can be fitted to general geometries and small details can be resolved.
- Unstructured grids lead to more elaborate memory accessing, and the number of operations per cell or wavelength is higher than in the classical FD-TD method.

Coupling of structured grids and unstructured grids strives to utilize their advantages without suffering too much from their disadvantages. For unstructured grids we distinguish between two cases:

1. The cell sizes in the unstructured domain are of the same size as in the Cartesian grid. Typically this is the case when moderately curved boundaries are meshed.

2. The cell sizes in the unstructured domain vary from the same size as in the Cartesian grid to very small cells. The small cells are used close to fine geometrical details in order to capture the physical influence of these details on the global solution.

We distinguish between two types of time integrators: explicit and implicit time integrators. The principal difference is:

$$\mathbf{u}^{n+1} = \mathbf{B}u^n + \mathbf{f} \quad (\text{explicit}), \quad (6.51)$$

$$\mathbf{K}u^{n+1} = \mathbf{B}u^n + \mathbf{f} \quad (\text{implicit}), \quad (6.52)$$

where \mathbf{u}^n is the vector of discrete unknowns at time level n , \mathbf{u}^{n+1} is the vector of discrete unknowns at time level $n + 1$, \mathbf{B} and \mathbf{K} are (non-diagonal) matrices (for finite elements a combination of the stiffness- and mass-matrix) and \mathbf{f} is a source term vector.

For explicit timestepping algorithms the time step is limited by the smallest cell size. This is generally not the case for implicit timestepping algorithms, where the time step usually can be chosen arbitrarily regarding stability. On the other hand, explicitly marching on in time is trivial since this is only a straightforward matrix-vector multiplication whereas implicitly marching on in time gives a system of linear equations to solve for each time step. However, the mass matrix in (6.52) is time independent and hence an LU factorization can be done in the initialization phase. But implicit algorithms cannot, in general, be used for large unstructured grids since the computational cost of doing an LU factorization does not scale linearly with the number of cells. An option to circumvent this is to use iterative methods to solve (6.52).

For large unstructured grids where some cells are small, explicit and implicit solvers will be very computer demanding if used separately. In that case one could try to couple an explicit solver (used for the larger cells) with an implicit solver (used for the smaller cells). This is outlined in Section 6.6.3.

There is an appropriate note to make here. Using implicit timestepping means that the time step does not have to be in accordance with the space step. Hence the high spatial frequencies that can be (locally) supported by regions of small cells cannot be propagated by the timestepping mechanism because larger time steps only support temporal frequencies up to a limit of

$$\max(f_{\text{supp}}) = \frac{1}{2 * \Delta t} \ll \frac{c}{2 * \Delta},$$

where Δ is the smallest cell size. A justified question is then why bother to resolve small details? The answer is that small geometric details do have an impact on the lower part of the (temporal) frequency spectrum part of the solution. This is easily realized by thinking of a thin infinite large PEC plate with a small hole. Even though the hole might be smaller than a wavelength, a wave impinging the plate on one side will result in fields excited into the other side.

6.6.1 FD-FE

Coupling between the FD-TD solver for structured grids and the implicit FE solver for unstructured grids (described in Section 6.3) is shown in Figure 6.12 and Figure 6.13.




Figure 6.12. Coupling between structured FD-TD and unstructured FE. The locations of the H_y components on Γ_A and Γ_a coincide. The locations of the E_z components on Γ_B and Γ_b also coincide.

1. The H_x and H_y components in the structured grid (to the left in Figure 6.12) are updated using the standard FD-TD method once E_z on Γ_B is given (Γ_B is the FD-TD boundary at B). When H_x and H_y are calculated the H_y components along Γ_A are sent to Γ_a in the FE grid.
2. The E_z components in the structured domain are updated using the standard FD-TD method.
3. The FE solver calculates the magnetic field implicitly in the unstructured domain using the values at Γ_a as boundary conditions.
4. By utilizing a discrete ∇ -operation around the auxiliary \tilde{E}_z variables, \tilde{E}_z can be time stepped according to $\tilde{E}_z^{n+1} = \tilde{E}_z^n + \frac{\Delta t}{\epsilon} \nabla \times \mathbf{H}$ at Γ_b .

With the grid we use at the interface the updating of \tilde{E}_z becomes identical to the standard FD-TD method and when \tilde{E}_z is known on Γ_b these variables are copied to Γ_B for the next time step.

6.6.2 FD-FV

Coupling between the FD-TD solver for structured grids and the FV solver for unstructured grids (described in Section 6.4) is very similar to the procedure described in the previous section. However, one main difference is important to point out.

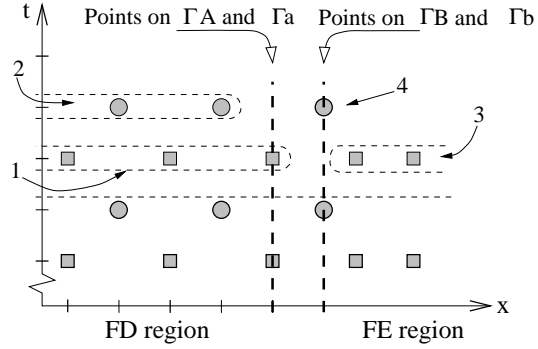


Figure 6.13. The timestepping mechanism for the FD-FE hybrid. Compare with Figure 6.1.

The FD-TD solver sends E_z components to the FV solver and receives magnetic components from the FV solver, see Figure 6.14.




Figure 6.14. Coupling between structured FD-TD and unstructured FV. The locations of the E_z components on Γ_A and Γ_a coincide. The locations of the H_y components on Γ_B and Γ_b also coincide. (On Γ_b , the H_y components are denoted $\tilde{\mathbf{H}} \cdot \mathbf{t}^p$ to signify that they are edge values.)

In Figure 6.14 only the magnetic edge values along Γ_b are drawn (denoted $\tilde{\mathbf{H}} \cdot \mathbf{t}^p$). They are collocated with the H_y components along Γ_B in the FD-TD region and hence a straightforward sending of magnetic components from the FV domain at Γ_b to Γ_B in the FD-TD domain can be performed.

6.6.3 FE-FV

A coupling between the two unstructured solvers is rendered possible because of the differences in sending to and receiving from the two unstructured solvers when cou-

pled with the structured FD-TD solver. The FE solver sends E_z components to the FV solver and receives magnetic components from the FV solver, see Figure 6.15.

In this case one can use the explicit solver for the larger cells and the implicit solver where smaller cells otherwise would limit the time step.




Figure 6.15. Coupling between unstructured FE and unstructured FV. The locations of the E_z components on Γ_A and Γ_a coincide. The locations of the H components on Γ_B and Γ_b also coincide. (On Γ_b , the H components are denoted $\tilde{H} \cdot t^p$ to signify that they are edge values.)

Our hybrid code is currently limited to coupling between FD-TD and one of the two unstructured solvers which means that the coupling of the two unstructured solvers has not been verified experimentally. There are no reasons to believe that there would be any accuracy problems but one cannot say anything about stability without numerical experiments.

6.7 Stability

A critical aspect of every numerical method is stability. All three schemes described in Sections 6.2–6.4 are stable as long as the CFL limits in (4.12) and (6.48) are not violated. However, this does not guarantee stability for the hybrid schemes since coupling of two stable schemes might result in an unstable scheme.

The complexity of the hybrid scheme makes it very difficult and cumbersome to perform a theoretical stability analysis. Instead we perform a numerical study. The entire hybrid scheme can be written as

$$\mathbf{u}^{n+1} = \mathbf{A}\mathbf{u}^n + \mathbf{f}^n, \quad (6.53)$$

where \mathbf{f}^n represent the source terms. The vector \mathbf{u} contains the unknown values of both \mathbf{E} and \mathbf{H} fields and includes unknowns from several time levels from both the structured domain and the unstructured domain. Furthermore \mathbf{u}^0 is given by initial values. Instability in the sense of exponential growth is generated if any eigenvalue

of \mathbf{A} is larger than one. One possible way to find the eigenvalues of \mathbf{A} would be to use the technique employed for the FV solver in Section 6.4.7. However, it would be very cumbersome to find the explicit values for all the elements of \mathbf{A} in the hybrid case. An alternative technique is to examine the dominant eigenvalue of \mathbf{A} by running the code for a very large number of time steps. We have utilized this procedure which in numerical linear algebra is known as the power method.

6.7.1 Details

We performed stability tests for several hybrid grids, including those three grids used in the convergence study for the circular cylinder, see Section 6.8. In all cases we tested both PEC and PMC boundary conditions.

We ran the code for ten million time steps. The time to complete these computations on an IBM Power 3 processor ranged from a few hours up to 24 hours depending on the size of the grid. As ABC we used a twelve cell thick U-PML layer with the profile given in (6.1). A plane wave was continuously fed into the computational domain using Huygens' surfaces. The actual shape of the pulse is not so important. It is crucial that all frequencies supported by the grid are present in the calculation, but round off errors will introduce them even if they are not present initially. A clever choice of excitation may lead to a quicker discovery of instability, but instability will finally show up no matter what pulse is used.

6.7.2 FD-FE

We did encounter stability problems for the FD-FE hybrid for some grids when the Crank-Nicholson scheme was used for time discretization. The grids used in Section 6.8 were all stable, with one exception. The finest grid was unstable when PMC BC was used. Furthermore, another grid, which is very similar to the coarse grid used for the circular cylinder, showed instabilities for both PEC and PMC.

Stability problems with FD-FE hybrids are not unique to our approach. It has been noted by other researchers in this area. A remedy has been suggested by Hwang and Wu [51]. They used a temporal filtering technique to stabilize their scheme. However, this approach reduces the order of convergence to one.

Our technique to stabilize the unstable cases is to increase the value of θ in the timestepping scheme, see (6.29). Thus, we increase the stability region and introduce dissipation for purely imaginary eigenvalues of the spatial discretization matrix. However, $\theta = 0.5$ is the only value that gives a second-order time integration scheme. Table 6.4 summarizes our results for the three grids that were unstable using the θ -method with $\theta = 0.5$, i.e. Crank-Nicholson.

θ	coarse grid, PEC	coarse grid, PMC	fine grid, PMC
0.50001	Unstable	Unstable	Unstable
0.5001	Unstable	Stable	Stable
0.501	Stable	Stable	Stable

Table 6.4. Result of stability tests for the FD-FE hybrid.

For all three grids it was enough to increase θ to 0.501. This change in θ is so small that the impact on accuracy is negligible even though it formally reduces the order of convergence (in time) to one.

We have not encountered any instabilities when using the BDF-2 timestepping scheme. The drawback of this scheme is that it has a time integration error approximately four times larger than the Crank-Nicholson scheme.

6.7.3 FD-FV

All grids tested were stable for the FD-FV hybrid as long as the time step was selected properly. The time step we used was a factor of $\sqrt{2}/4$ smaller than the stability limit for the FD-TD scheme, i.e. we chose

$$\Delta t = \frac{\sqrt{2}}{4c_0 \sqrt{\frac{1}{\Delta x^2} + \frac{1}{\Delta y^2}}}. \quad (6.54)$$

Using a time step twice as large proved to be unstable. This is not surprising since it violates the stability condition in (6.48). The grids used in these stability tests all have a shortest edge length of slightly less than half the cell size of the structured grid.

6.8 Convergence

The order of convergence is at least two in both time and space for all three schemes described in Sections 6.2–6.4. The hybridization techniques presented in Section 6.6 are designed to preserve this property. We have validated this by doing calculations on five different cases:

- vacuum,
- a circular PMC cylinder,
- a circular PEC cylinder,
- a circular cylinder with $\epsilon_r = 4$, and
- a circular cylinder with $\mu_r = 4$.

Note that the duality of the Maxwell equations means that PMC boundary condition for the TM mode is equivalent to PEC boundary condition for the TE mode. Hence a PMC cylinder is an interesting test case even though there are no PMC materials in the real world. We compare with FD-TD solutions which use staircase approximations of the circular cylinder. We demonstrate that when the circular cylinder is present, the FD-TD method does not show second-order convergence but the hybrid methods do. This holds for all four test cases with a circular cylinder. In all cases, we have used Huygens' surfaces to create a plane wave, with a Gaussian shape given by

$$f(t) = e^{-(t-t_0)^2/t_w^2}, \quad (6.55)$$

with $t_0 = 20/c_0$ and $t_w = t_0/6$. The angle of incidence was given by $\mathbf{k} = (1, 0)$. As ABC we used a twelve cell thick U-PML layer with the profile given in (6.1).

6.8.1 Modeling of circular cylinders in FD-TD

As explained in Section 4.6, we have to model the circular cylinder using a staircased approximation when we use the FD-TD method. The PEC and PMC boundary conditions are implemented by zeroing components on the boundary of the staircase representation of the circular cylinder. However, we still have to choose an algorithm for finding this representation of the circular cylinder.

In the PEC case, a simple way would be to zero all E_z components that are located inside the circular cylinder. However, this would lead to a staircase representation that always has an area smaller than the true area of the circular cylinder. Instead, we have chosen to zero all E_z components that are located inside a virtual radius of the true radius plus half the cell size. (We use the same resolution in both spatial dimensions.) Numerical tests have confirmed our belief that this procedure gives better results.

PMC boundary condition is implemented by zeroing the four surrounding \mathbf{H} components of every E_z component that are located inside the circular cylinder, see Figure 4.1.

A computational cell is considered to be a part of the cylinder if its center lies inside the cylinder. An E_z component is located at the corners of four cells, see Figure 4.1. The ϵ value needed for the update of this E_z component is taken as the arithmetic mean value of the ϵ values in these four cells. The H_x and H_y components are located on the sides of the cells. The μ values needed for the update of these components are taken as the harmonic mean value of μ in the two cells sharing this side. The reason for using a harmonic mean value rather than an arithmetic is the fact that these components are normal to the interface. This issue is discussed in detail in the Ph. D. thesis of Ulf Andersson [5].

6.8.2 The coarse grids

Figure 6.16 displays the unstructured grids for a circular cylinder with a radius of two meters. These 8×8 meter grids were created using Femlab [35] and inserted




Figure 6.16. The coarsest unstructured grids.

in the center of a structured FD-TD grid with 40×40 cells of the size $\Delta = 1$ m so that the center of the cylinder was located at $(x,y)=(19.5, 20.5)$. The coarsest unstructured grid for the vacuum calculations consists of eight by eight squares that have been split into two triangles each.

In all FD-TD calculations we have used $CFL = 1/\sqrt{2}$ for the circular cylinder. The same time step was used for the FD-FE hybrid, while the time step for the FD-FV hybrid was a factor of two smaller for the PEC and PMC cylinders and a factor of four smaller for the material cylinders. For $\Delta x = \Delta y = \Delta = 1$ m this gives $\Delta t \approx 1.67$ ns, $\Delta t \approx 0.83$ ns and $\Delta t \approx 0.42$ ns. The FD-FV time step has not been chosen as the largest possible value. It has deliberately been chosen so that the FD-FE time step is a multiple of the FD-FV time step. This allows us to avoid temporal interpolation when probing the solutions. In the vacuum case all schemes have used $CFL = 1/\sqrt{8}$. We have used Crank-Nicholson for the FE timestepping.

The error has been measured in those E_z components in the FD-TD grid that lie closest to the unstructured grid. There are 36 such components, namely the outer components in the set $E_z(17:24,18:25)$. By choosing to measure the error in the FD-TD grid we avoid spatial interpolation since all refined grids will have E_z components in these locations. The error is defined as the difference between the computational solution and a reference solution. The reference solutions for the circular cylinders have been obtained by highly resolved FD-TD calculations. We used $N_x = N_y = 10240$ for the PEC and PMC cylinders and $N_x = N_y = 5120$ for the material cylinders. The errors in these reference solutions have been estimated by comparing them with solutions where the problem size was a factor of two smaller in each dimension. The results are given in Table 6.5.

	Error estimate
Vacuum	0
PEC	$1.3 \cdot 10^{-4}$
PMC	$1.6 \cdot 10^{-4}$
$\epsilon_r = 4$	$2.8 \cdot 10^{-5}$
$\mu_r = 4$	$3.8 \cdot 10^{-4}$

Table 6.5. Estimates of the maximum errors in the reference solutions.

For vacuum, we have used the analytical solution as the reference solution.

The Huygens' surfaces were placed two meters from the outer boundary. In this case it means that they were placed in the third cell. When we refine the grid we will keep the physical location of the Huygens' surfaces at two meters from the boundary.

6.8.3 Grid refinement

The unstructured grids in Figure 6.16 consist of two parts: the interior 6×6 meter part and a one cell thick transition layer. When we refine the grids, the interior part of the unstructured grids are refined by splitting all triangles into four triangles. This is done in such a way that all edges are split into two edges. This was done using a slight adjustment in the grid refinement procedure of **Femlab** [35]. When splitting one right angle triangle into four new triangles, **Femlab** splits the triangles as in the left part of Figure 6.17. As mentioned in Section 6.5 we prefer the methodology illustrated in the right part of Figure 6.17.




Figure 6.17. Two different grid refinement methodologies.

After the refinement of the interior part, we add a one cell thick transition layer outside the interior part. This means that the area of the unstructured grid shrinks with the refinement, see Table 6.6.

The nature of the hybrid scheme and the refinement procedure described above make it necessary to move the circular cylinder when a grid is refined with a factor of two. The lower left corner of the interior part of the unstructured grid must coincide with the center of a cell in the structured grid, i.e. a location where there is no field component. When we refine the grid with a factor of two, a new E_z component will appear in the center of the coarse grid cell and hence, we may not have the lower left corner of the interior part there. Hence, we have to move the unstructured grid. The location of the center of the circular cylinder, $(xcyl, ycyl)$,

is specified in Table 6.6. A separate reference solution was calculated for each location of the circular cylinder. The relocation of the cylinder could have been avoided by using a refinement factor of three, but we considered this to be a too rapid refinement.

Grid	Coarse	Medium	Fine
$N_x = N_y$	40	80	160
Δ (m)	1	0.5	0.25
side of unstr. domain (m)	8	7	6.5
PEC, area (m ²) (FD-TD)	9	11.25	11.56
PMC, area (m ²) (FD-TD)	12	13	13
xcyl (m)	19.5	19.25	19.125
ycyl (m)	20.5	20.25	20.125
number of triangles	100	280	904
number of edges	172	460	1432
shortest dual edge	0.4714	0.2357	0.1179
Bandwidth (FE)	14	19	36

Table 6.6. Parameters for the circular cylinder test case. (PEC and PMC)

Grid refinement of the semi-structured grid used for the vacuum test case was performed in a similar manner. However, in this case we need only one reference solution since there is no object in the unstructured grid.

6.8.4 Results

The results for the convergence study are shown in Figures 6.18–6.22. The plots display the mean value of the absolute error in the 36 measure points as a function of time. The peak of the Gaussian pulse passes the circular cylinder at $t \approx 125$ ns.




Figure 6.18. Errors for three refinement levels and three methods for the circular PMC cylinder.




Figure 6.19. Errors for three refinement levels and three methods for the circular PEC cylinder.

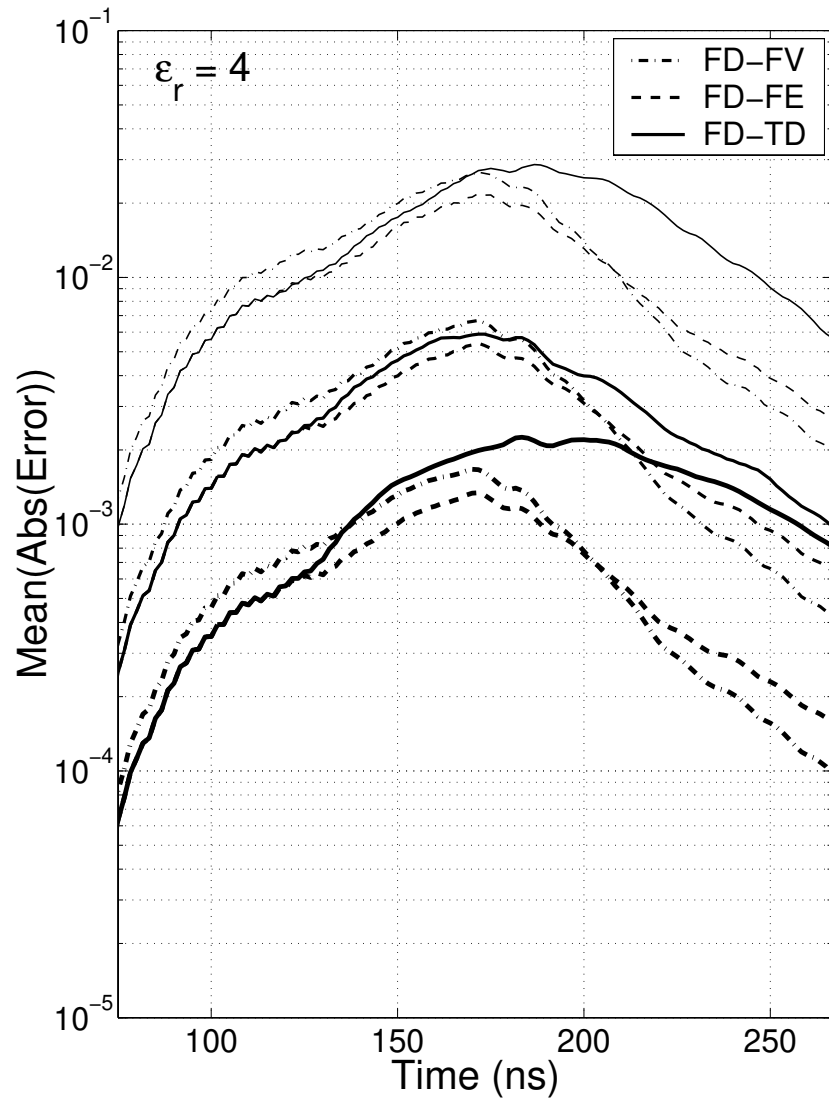


Figure 6.20. Errors for three refinement levels and three methods for the circular cylinder with $\epsilon_r = 4$.




Figure 6.21. Errors for three refinement levels and three methods for the circular cylinder with $\mu_r = 4$.




Figure 6.22. Errors for three refinement levels and three methods for vacuum.

We note that in the beginning, the error is larger for the FD-FV hybrid. This is caused by the shorter time step which gives a larger error in the FD-TD scheme. This effect is not present in Figure 6.22 because we have used the same time step for all methods in this particular test case.

For PEC and PMC, we note that around 110–115 ns, the error in the FD-TD solution is slightly smaller than the errors in the hybrid solutions. This is most likely due to reflections in the transition layer between the structured grid and the unstructured grid. However, once the errors in the geometrical representation of the circular cylinder affects the solution, we clearly see that the hybrid schemes are superior.

The increase in error that can be seen in Figure 6.22 for $t > 220$ ns is caused by the Huygens' surfaces. When the Gaussian pulse reaches the upper Huygens' surface a Gaussian with opposite sign is generated which is intended to zero out the approaching Gaussian and thus giving a zero scattered field. Due to the dispersion error there is a mismatch leading to a nonzero scattered field. This error is however absorbed by the ABC. The error mentioned above is an error component generated at the upper Huygens' surface and traveling in the opposite direction. This effect is present also in the other four cases, but is harder to spot since errors in the scattered field are also present.

The order of convergence has been estimated by calculating the L_1 norm of the errors in Figures 6.18–6.22. The result is given in Table 6.7. In all cases we have second-order convergence for the hybrid schemes. The FD-TD scheme however, only has second-order convergence for vacuum. Note that we have used a numerical reference solution for the circular cylinder cases. This affects the convergence estimates in Table 6.7. The last digit could be slightly altered if we use a better reference solution. However, this does not alter the main conclusion that we have second-order convergence for the hybrid schemes.

We would also like to point out that the results in Figures 6.18–6.22 prove that there are no large reflections at the grid interfaces. Small reflections are unavoidable, but as long as these reflection errors do not dominate the other error sources, they are acceptable.

	FD-TD		FD-FE		FD-FV	
	coar	fine	coar	fine	coar	fine
Vacuum	2.08	2.02	2.08	2.02	2.06	2.02
PEC	2.23	1.06	2.01	1.98	1.99	1.98
PMC	1.33	0.97	2.13	2.08	2.01	1.91
$\epsilon_r = 4$	1.50	1.10	2.03	2.00	2.03	2.00
$\mu_r = 4$	1.54	1.16	2.06	2.01	2.04	1.97

Table 6.7. Estimates of the order of convergence for the three methods. The values in the columns labeled coar have been obtained by comparing errors on the coarse and medium grids and the values in the columns labeled fine have been obtained by comparing errors on the medium and fine grids.

6.9 PMC wall

A classical paper demonstrating the errors caused by staircasing is that by Cangellaris and Wright [17]. They study waves propagating along PEC walls in 2D and conclude that the TE modes suffer dispersion due to the staircase approximation of a PEC wall, while the TM modes do not.

One of their test cases involved a line source close to a PEC wall. The field was probed at a distance of 100 cells from the line source. We make a similar test here. Since we are working with the TM equations, we use the PMC boundary condition. The duality of the Maxwell equations makes this equivalent to the PEC boundary condition for the TE equations, which was shown by Cangellaris and Wright [17] to be the “bad” case.

6.9.1 Details of the numerical setup

The computational domain is a square, i.e. we have $N = N_x = N_y$ and $\Delta = \Delta x = \Delta y$. The value of N is chosen such that effects from the ABC or from the edges of the PMC wall do not reach the probing points during our probing window, i.e. N depends on the number of time steps taken. We use $\Delta = 0.1$ m and $CFL = 0.75$. This gives us $\Delta t \approx 0.176$ ns.

The line source is applied as a current source, i.e.

$$E_z|_{i,j}^{n+1} = E_z|_{i,j}^{n+1} + \frac{\Delta t}{\epsilon \Delta^2} e^{-((n+\frac{1}{2})\Delta t - t_0)^2 / t_w^2}, \quad (6.56)$$

where $t_0 = 10.56$ ns and $t_w = t_0/6 = 1.76$ ns. The weighting with $\Delta t/\Delta^2$ is necessary to ensure convergence since the discrete current source is equivalent to an integral over the cell area which is Δ^2 . The Δt comes from the time discretization.

In the same way as in [17] we made two different FD-TD calculations: one where the PMC wall aligns with the FD-TD grid and one where the wall is tilted 45 degrees compared to the grid axes. We will refer to the first case as “para” and the second as “diag”.

In the “para” test case we placed the line source at $E_z(N/2-50,8)$ and the PMC wall at $H_x(:,2)$. Hence, the distance from the wall to the line source was $h = 0.65$ m. The probing point was placed in $E_z(N/2+50,8)$, i.e. it lies at the same distance from the wall as the line source. The distance between the line source and the probe point were $d = 10.0$ m.

In the “diag” test case we placed the line source at $E_z(N/2-54,N/2-46)$ and the PMC wall was modeled by zeroing $H_y(k,k)$, $k=1,\dots,N$, and $H_x(k+1,k)$, $k=1,\dots,N-1$ in each time step. If we regard the wall as being placed along the line $y=x-\Delta/2$, then the distance from the wall to the line source is $h = 0.60$ m. The probing point was located at $E_z(N/2+17,N/2+25)$. Hence, the line source and the probe point were separated by a distance of $d = 10.04$ m.

To obliterate the error introduced by the staircase approximation of the tilted wall, we introduced an unstructured grid close to the wall, see Figure 6.23. The line

source and the probing point are located as in the “diag” case. The distance h from the PMC wall to the source point and probe point was $h = 4.5 \cdot \sqrt{2}\Delta \approx 0.636$ m. Due to the one half cell overlap between the structured and unstructured grid, it is not possible to align the wall along $y=x-\Delta/2$. In Figure 6.23 we only display a small piece of the unstructured grid. In all calculations we have used a large enough grid to insure that effects from the edges of the PMC wall do not influence the solution at our probing point.




Figure 6.23. The horizontal wall is shown to the left and the tilted wall in the middle. The unstructured grid used for the tilted wall is shown to the right.

The FV hybrid was not stable for $CFL = 0.75$. Instead we used $CFL = 0.375$, doubled the number of time steps and probed the field every other time step.

6.9.2 Results

Figure 6.24 presents the result of this test case. The analytical solution is given by

$$u(t) = \sum_{i=1}^2 \int_{-\infty}^t q(\tau) G(\rho_i, t - \tau) d\tau, \quad (6.57)$$

where ρ_1 is the distance between the current source and the probing point, ρ_2 is the distance between the image point of the current source and the probing point,

$$q(\tau) = \frac{1}{\epsilon_0} \frac{\partial}{\partial \tau} e^{-(\tau-t_0)^2/t_w^2}, \quad (6.58)$$

and the Greens function is given by

$$G(\rho, t) = \frac{H(tc_0 - \rho)}{2\pi c_0 \sqrt{t^2 - \rho^2/c_0^2}}, \quad (6.59)$$

where $H(t)$ is the Heaviside step function. Note that ρ_1 and ρ_2 are slightly different for “para”, “diag” and the unstructured solutions. Hence we calculate different analytical solutions for these cases. The analytical solution displayed in Figure 6.24 is for the “para” case.

Comparing the errors in Figure 6.24 we clearly see that the hybrid schemes outperform the “diag” case. We also note that the errors of the two hybrid schemes are very similar. Furthermore, the hybrid solutions are actually slightly better than the “para” solution. This may be due to the fact that the error in the FD-TD scheme is smaller for waves traveling diagonally than for waves traveling along a Cartesian axis, see Chapter 5 in [98].

We also tested two contour path modeling schemes. The first scheme is the one described in Chapter 10.6 of [98]. It is labeled “Taf” in Figure 6.24. This scheme was found to be unstable, but the instability did not significantly affect the solution in the probing point during our probing window. The instability was generated at the wall. Finally, we also implemented the scheme of Dey and Mittra [21]. This scheme was found to be stable for this test case.




Figure 6.24. The analytical solution in the probing point and the errors of the six schemes.

We note that these contour path modeling schemes give much better results than the “diag” solution, but they are not quite as good as our hybrid schemes. Furthermore, even though the contour path modeling schemes were rather easy to implement for this very special 2D case, we feel that it would be a very complex and cumbersome task to make a general implementation in 3D.

6.10 Conclusion

We have introduced a new general hybrid approach for solving the Maxwell equations in the time domain. By combining the efficiency of the classical FD-TD method with the flexibility of solvers for unstructured grids we obtain a very favorable compromise between efficiency and accuracy. Flexibility is further enhanced by using two solvers on the unstructured grids, one explicit FV solver and one implicit FE solver. Note that the key words here are implicit and explicit and not FE and FV. An explicit solver is much faster per time step than an implicit solver. On the other hand, we have unconditional stability for the implicit solver while the explicit solver must obey a stability limit where the maximum time step is proportional to the shortest edge in the unstructured grid. Furthermore, the accuracy in the FD-TD scheme decreases when the time step decreases, since the spatial and temporal errors are of opposite sign. Hence, the proper choice of solver for an unstructured part in the hybrid depends on the cell sizes in the unstructured grid.

We chose not to put any extra effort into optimizing the unstructured 2D solvers. Hence, we feel that it might be misleading to make any measurements of the efficiency in 2D. Furthermore, we feel that this question is much more relevant in 3D.

Chapter 7

Hybrid Methods in 3D

7.1 Introduction

This chapter is devoted to the 3D hybrid solvers. The implicit Finite-element solver is described in Section 7.3. An adoption of the plane wave excitation originally developed for FD-TD is also included. It is primarily intended for wave guide excitation in the FE solver. In Section 7.4 a thorough description of the explicit Finite-volume solver is given. It also contains a divergence analysis as well as a stability analysis for the FV solver. This section is based on the Licentiate thesis of Fredrik Edelvik [29]. Section 7.5 describes the hybridization technique used to couple the structured and unstructured solvers together. The optimal choice of method on the unstructured grid depends on the cell sizes. An implicit method is preferable for grids with small cells while an explicit method is preferable for grids where the cell sizes are of the same order as in the Cartesian grid. The stability of the hybridization is discussed in Section 7.6. The stability is not absolute and instabilities are observed for highly resonant geometries. But for all scattering problems examined the stability has not been an issue. In Section 7.7 several results are presented. Convergence results on simple vacuum grids shows that the solvers are second-order accurate but under hybridization we formally lose accuracy. However, the accuracy in the scattering problems examined are very satisfactory for the hybrid solvers. Finally, conclusions are given in Section 7.8.

7.2 Finite-difference method

The FD-TD method is described in Chapter 4 and will not be further discussed here. We simply note that we use the uniaxial perfectly matched layer (U-PML) absorbing boundary conditions [38], and for plane wave excitation the total/scattered field decomposition described in [98] is used. The FD-TD method is also thoroughly described in the literature, see for example the book by Taflovie [98].

7.3 Finite-element method

In this section the Finite-Element Time-Domain (FE-TD) method in 3D is briefly discussed. A lot of details are the same as in 2D and hence Section 6.3.1 contains materials also valid here.

In 3D we use the vector wave equation for the electric field, instead of the magnetic field formulation used in 2D TM case. For $\sigma = \sigma^* = 0$ we have

$$\epsilon \frac{\partial^2 \mathbf{E}}{\partial t^2} + \nabla \times \frac{1}{\mu} \nabla \times \mathbf{E} = \mathbf{0}. \quad (7.1)$$

The weak formulation is now: find $\mathbf{E} \in W$ such that

$$\int_{\Omega} \left(\epsilon \frac{\partial^2 \mathbf{E}}{\partial t^2} \cdot \mathbf{w} + \frac{1}{\mu} \nabla \times \mathbf{E} \cdot \nabla \times \mathbf{w} \right) d\Omega = - \int_{\Gamma} \frac{1}{\mu} \mathbf{n} \times \nabla \times \mathbf{E} \cdot \mathbf{w} ds, \quad (7.2)$$

for all $\mathbf{w} \in W$. Again we use “edge” or “Whitney” elements but they are now tetrahedral elements [76] instead of triangular.

To define these linear edge elements, which are second-order accurate, consider the standard linear basis functions Φ_i for nodal-based finite elements, constructed such that $\Phi_i = 1$ in node number i and $\Phi_i = 0$ in all other nodes. Take edge e to be the edge on a tetrahedral element joining node i and node j . The basis function for edge e is

$$\varphi_e = \Phi_i \nabla \Phi_j - \Phi_j \nabla \Phi_i. \quad (7.3)$$

The properties of the basis functions are the same as in 2D,

- $\nabla \cdot \varphi_e = 0$.
- φ_e has constant tangential component (= 1/length) along edge e , which means that the tangential component is continuous around edge e . The normal components are discontinuous across element boundaries for φ_e .
- φ_e has zero tangential component along the other edges.

A sketch of the basis function for edge e is plotted in Figure 7.1.

Our unknowns can be seen as edge values because there is only one basis function that has a nonzero tangential component on each edge. Furthermore, we can always use the basis functions to calculate the electric field in any point inside each tetrahedra. But if we seek the electric field at the nodes we are in trouble because the basis functions are not continuous around a node. In this case a least square fit of all emanating edges can be used.




Figure 7.1. The vector valued basis function φ_e for the edge connecting node 1 and 2. φ_e is illustrated for this tetrahedral element but is also defined in all other tetrahedra that share this edge. (Node 3 in the tetrahedron is furthest away.)

Since we are solving for the \mathbf{E} field we use a Neumann boundary condition for PMC, a homogeneous Dirichlet boundary condition for PEC and an inhomogeneous Dirichlet boundary condition for the hybrid boundary,

$$\mathbf{n} \times (\nabla \times \mathbf{E}) = \mathbf{0} \quad \text{on } \Gamma_{\text{PMC}}, \quad (7.4)$$

$$\mathbf{n} \times \mathbf{E} = \mathbf{0} \quad \text{on } \Gamma_{\text{PEC}}, \quad (7.5)$$

$$\mathbf{n} \times \mathbf{E} = \mathbf{n} \times \mathbf{E}_{\text{fdd}}(t) \quad \text{on } \Gamma_{\text{Hyb}}. \quad (7.6)$$

We use the same timestepping technique as in 2D, see Section 6.3.3 for a description. We also use the same procedure to renumber and solve the system of linear equations. The use of a direct method to solve the linear system of equations is a major bottleneck for larger problems in 3D. Therefore we have recently implemented an iterative solver based on the conjugate gradient method. An incomplete Cholesky factorization is also used as preconditioner. Preliminary results shows that the iterative method converges within tens of iterations. However, this is not further discussed here.

7.3.1 Huygens' surfaces for the FE solver

Traditionally Huygens' surfaces are associated with plane wave excitations in the FD-TD method. But the same idea can also be utilized and formulated for the FE-TD method. Furthermore, as in FD-TD, Huygens' surfaces can also be used for waveguide excitation. In fact, Huygens' surfaces can be used for any excitation that can be formulated over a surface, for example plane wave expansions, waveguide modes and spherical expansions. The only requirement on the excitation is that the electric field is known for each edge, as a function of time, around the Huygens' surface.

The Huygen excitation is based on splitting the total field in a scattered and an incident field, i.e.

$$\mathbf{E}^{\text{tot}} = \mathbf{E}^{\text{sca}} + \mathbf{E}^{\text{inc}}. \quad (7.7)$$

Let us start the derivation by the following formulation of the spatial discretization,

$$M\ddot{\mathbf{E}} + S\mathbf{E} = 0, \quad (7.8)$$

where conductivity and boundary terms are neglected for simplicity. Furthermore, split the matrices in 3×3 blocks and the unknowns in three parts, which yields

$$\begin{bmatrix} M_{11} & M_{12} & \mathbf{0} \\ M_{21} & M_{22} & M_{23} \\ \mathbf{0} & M_{32} & M_{33} \end{bmatrix} \begin{bmatrix} \ddot{E}_1^{\text{tot}} \\ \ddot{E}_2^{\text{tot}} \\ \ddot{E}_3^{\text{tot}} \end{bmatrix} + \begin{bmatrix} S_{11} & S_{12} & \mathbf{0} \\ S_{21} & S_{22} & S_{23} \\ \mathbf{0} & S_{32} & S_{33} \end{bmatrix} \begin{bmatrix} E_1^{\text{tot}} \\ E_2^{\text{tot}} \\ E_3^{\text{tot}} \end{bmatrix} = 0, \quad (7.9)$$

where E_1^{tot} , E_2^{tot} , and E_3^{tot} contains three groups of unknowns which are written in a total field formulation. Now we can substitute the total field by the scattered field using (7.7) because the only difference is the known incident field. Therefore, let us associate the edges in the first region to a scattered field formulation, and the edges in the third region to a total field formulation. Furthermore, associate the edges on the triangulated surface between the two regions to a scattered field formulation. The surface is called a Huygens' surface and the corresponding edges forms the second group of unknowns. This is illustrated in Figure 7.2 where two tetrahedral elements are separated by a Huygens' surface.




Figure 7.2. Scattered field formulated unknowns in region 1, and total field formulated unknowns in region 3 are separated by the Huygens' surface. The edges on the surface is associated to scattered field formulated unknowns.

The edges in the first and third region are not directly coupled to each other because the basis functions have a local support and the common edges are associated to the Huygens' surface. Therefore we have two zero-blocks in the block-matrix structure of (7.9).

Consider the first and second group of unknowns to be solved only in scattered field formulation. Split E_1^{tot} and E_2^{tot} into the corresponding scattered and incident variables using (7.7). This gives

$$\begin{bmatrix} M_{11} & M_{12} & \mathbf{0} \\ M_{21} & M_{22} & M_{23} \\ \mathbf{0} & M_{32} & M_{33} \end{bmatrix} \begin{bmatrix} \ddot{E}_1^{\text{sca}} \\ \ddot{E}_2^{\text{sca}} \\ \ddot{E}_3^{\text{tot}} \end{bmatrix} + \dots = \begin{bmatrix} \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & -M_{23} \\ \mathbf{0} & M_{32} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \ddot{E}_1^{\text{inc}} \\ \ddot{E}_2^{\text{inc}} \\ \ddot{E}_3^{\text{inc}} \end{bmatrix} + \dots \quad (7.10)$$

where only the first block-matrixes are given. The other terms have identical structures.

If we take a closer look at where the incident field is applied we realize that the Huygens' surface is not mathematically collocated with the triangular surface. The surface divides the unknowns into the scattered and total field formulated unknowns and hence the surface is localized slightly to the right in Figure 7.2. However, it is not important to specify the position of the Huygens' surface exactly. But it is important to realize that the surface affects the edges in a layer around the surface. This layer is the key to the success of the Huygens' surface because it is not enough to specify the electric field on a surface to obtain an excitation propagating only in one direction.

The mechanism to excite the FE-TD solver is only described briefly and in the case of a plane wave excitation we simply use the analytically known incident field. This is verified numerically and results comparable to FD-TD excitations have been obtained. Also, if the unstructured grid can be characterized in terms of a dispersion relation the method of dispersion compensation described in [65] can be utilized.

Furthermore, we can use the Huygens' surface to excite wave guide modes and in this case we must find the time dependent incident mode solution for each edge affected by the surface. This is done in three steps; First, find the mode solution in frequency domain for the wave guide cross section. Next, evaluate the mode solution in the affected edges in the total field region. Third, Fourier transform the mode solution to time domain and excite the Huygens' surface with the time dependent incident field. This three step process can be very memory consuming if the wave guide is inhomogeneous because then all edges must have a unique time history. In the case of homogeneous cross sections the mode solution is frequency independent and each edge then only need an individual time dependent amplitude factor for each mode excited. However, this three step procedure is not further addressed in this thesis. It is currently being implemented in the GEMS project [39] and preliminary results are promising.

7.4 Finite-volume method

The FV solver is based on the following integral formulation of Faraday's and Ampère's laws given in (3.11):

$$\frac{\partial}{\partial t} \int_A \mu \mathbf{H} \cdot \mathbf{n} dA = - \oint_{\Gamma} \mathbf{E} \cdot d\mathbf{l}, \quad (7.11)$$

$$\frac{\partial}{\partial t} \int_A \epsilon \mathbf{E} \cdot \mathbf{n} dA = \oint_{\Gamma} \mathbf{H} \cdot d\mathbf{l} - \int_A \sigma \mathbf{E} \cdot \mathbf{n} dA, \quad (7.12)$$

where A is an arbitrary area, Γ is the path that encloses A and \mathbf{n} is a unit normal.

7.4.1 Space discretization

The integral formulations (7.11) and (7.12) are discretized on a staggered grid. The primary grid consists of tetrahedra, which are generated by a Delaunay grid algorithm. Associated with the unstructured tetrahedral grid is a mutually orthogonal dual grid, the Dirichlet tessellation, which is constructed in the preprocessing phase. The nodes in the dual grid are the centers of the circumscribed spheres of each tetrahedron. However, for a general tetrahedral grid there are circumcenters that lie outside the corresponding tetrahedra. Such nodes deteriorate the accuracy of the solver and they are hence replaced by the barycenters of these tetrahedra. Replacing a circumcenter with a barycenter implies that the mutual orthogonality is lost locally. The dual edges are constructed by joining dual nodes of tetrahedra sharing a common face. The dual face associated with a primary edge is defined by the dual edges which correspond to all the faces which share the given primary edge. The areas and normals are calculated by a piecewise planar approximation of the dual faces. Finally, the dual cell associated with a primary node is defined by the dual faces which correspond to all the primary edges sharing the given primary node. Hence, there is a one-to-one correspondence between the nodes, edges, faces and cells of the primary and dual grids. At each triangular face the normal component of the magnetic field is stored and, similarly, at each dual face the normal component of the electric field is stored, see Figure 7.3.

Integrating (7.11) and (7.12) over a primary and a dual face, respectively, results in

$$\tilde{\mu}_i^p \frac{\partial}{\partial t} \int_{A_i^p} \mathbf{H} \cdot \mathbf{n}_i^p dA = - \sum_m \int_{\Gamma_{i,m}^p} \tilde{\mathbf{E}} \cdot \mathbf{t}_{i,m}^p dl, \quad (7.13)$$

$$\tilde{\epsilon}_j^d \frac{\partial}{\partial t} \int_{A_j^d} \mathbf{E} \cdot \mathbf{n}_j^d dA = \sum_k \int_{\Gamma_{j,k}^d} \tilde{\mathbf{H}} \cdot \mathbf{t}_{j,k}^d dl - \tilde{\sigma}_j^d \int_{A_j^d} \mathbf{E} \cdot \mathbf{n}_j^d dA, \quad (7.14)$$

where A_i^p is the area of the primary face i , Γ_i^p is the path that encloses A_i^p , \mathbf{n}_i^p is the unit face normal for face i and $\mathbf{t}_{i,m}^p$ are the unit vectors for the edges in face i . The variables belonging to the dual face j are defined in the same way.




Figure 7.3. A cell in the primary grid and a dual face.

When the primary and dual grids are mutually orthogonal the normals and edge vectors are parallel. In that case the solver is basically a generalization of FD-TD to unstructured grids. The normal component of the electric field, $\mathbf{E} \cdot \mathbf{n}_j^d$, is updated by circulating the magnetic field along the edges of the dual face. The normal component of the magnetic field, $\mathbf{H} \cdot \mathbf{n}_i^p$, at each triangular face is updated in the dual way by circulating the electric field along the edges of the triangular face.

However, for non-orthogonal grids the normals and edge vectors do not align and we need an approximation for the electric and magnetic fields projected along the respective edges. This is accomplished by first calculating an approximate value of the fields at the primary and dual nodes, respectively. The magnetic field at the dual nodes, e.g. \mathbf{H}_r in Figure 7.3, are calculated as a least square fit of the magnetic fields normal to the four faces of each tetrahedron. Electric node values are defined by doing a similar fit for the electric fields normal to the dual faces of each dual cell. Following Riley and Turner [84], the magnetic field projected along the dual edge $\mathbf{t}_{j,k}^d$, see Figure 7.3, is then evaluated as

$$\tilde{\mathbf{H}} \cdot \mathbf{t}_{j,k}^d = (\mathbf{H} \cdot \mathbf{n}_i^p) (\mathbf{n}_i^p \cdot \mathbf{t}_{j,k}^d) + \frac{1}{2} [(\mathbf{H}_r + \mathbf{H}_q) - ((\mathbf{H}_r + \mathbf{H}_q) \cdot \mathbf{n}_i^p) \mathbf{n}_i^p] \cdot \mathbf{t}_{j,k}^d. \quad (7.15)$$

The electric field projected along the primary edge $\mathbf{t}_{i,m}^p$ is evaluated accordingly as

$$\tilde{\mathbf{E}} \cdot \mathbf{t}_{i,m}^p = (\mathbf{E} \cdot \mathbf{n}_j^d) (\mathbf{n}_j^d \cdot \mathbf{t}_{i,m}^p) + \frac{1}{2} [(\mathbf{E}_1 + \mathbf{E}_2) - ((\mathbf{E}_1 + \mathbf{E}_2) \cdot \mathbf{n}_j^d) \mathbf{n}_j^d] \cdot \mathbf{t}_{i,m}^p. \quad (7.16)$$

Taking a closer look at (7.15) and (7.16) we note that if \mathbf{n}_i^p is parallel to $\mathbf{t}_{j,k}^d$ and \mathbf{n}_j^d is parallel to $\mathbf{t}_{i,m}^p$, the second part of (7.15) and (7.16) vanish. Evaluating the fields projected along the edges in this way guarantees that the divergence of the electric and magnetic fields are preserved on a local cell level up to machine precision; see Section 7.4.3 for a proof. This has been found crucial in order to suppress spurious modes in the solution.

All materials are defined relative to primary grid cells. Hence, the material parameters require averaging to get an approximate value at the faces. The average permeability on primary face i , $\tilde{\mu}_i^p$, is defined by

$$\tilde{\mu}_i^p = \frac{2}{\frac{1}{\mu_r^p} + \frac{1}{\mu_q^p}}, \quad (7.17)$$

where r and q are the tetrahedra sharing face i . At a boundary face $\tilde{\mu}_i^p$ equals the permeability of the only tetrahedron associated with the face. The average electric material parameters on dual face j , $\tilde{\epsilon}_j^d$ and $\tilde{\sigma}_j^d$, are constructed as

$$\tilde{\epsilon}_j^d = \frac{\sum_q \epsilon_q^p A_{j,q}^d}{A_j^d}, \quad \tilde{\sigma}_j^d = \frac{\sum_q \sigma_q^p A_{j,q}^d}{A_j^d}, \quad (7.18)$$

where $A_{j,q}^d$ is the part of the area of the dual face which is inside tetrahedron q and the sums are taken over the tetrahedra associated with the dual face. Performing the material averaging in this manner preserves the second-order accuracy of the solver for inhomogeneous materials [5]. The reason for the different averaging of magnetic and electric materials is that the normal magnetic field is discontinuous at a material interface, whereas the tangential electric field is continuous. Therefore, harmonic averaging is used for μ and area weighted arithmetic averaging is used for ϵ and σ .

The boundary condition for a perfect electric conductor (PEC) is implemented as described in 6.4.2. For a perfect magnetic conductor (PMC) the tangential magnetic components, should equal zero at the boundary. However, the \mathbf{H} components are not situated on the boundary. The boundary condition is enforced by setting the magnetic field below the conductor equal to the value of the magnetic field directly above the conductor with a change of sign in the tangential component. The PMC boundary condition is non-physical, but useful as a symmetry boundary condition.

7.4.2 Time discretization

The equations (7.13) and (7.14) are discretized in time using a third-order staggered Adams–Bashforth scheme (ABS3) [40]

$$\begin{aligned} \mathbf{H} \cdot \mathbf{n}_i^p |^{n+\frac{1}{2}} &= \mathbf{H} \cdot \mathbf{n}_i^p |^{n-\frac{1}{2}} \\ &+ \frac{\Delta t}{\tilde{\mu}_i^p A_i^p} \sum_m \left(\frac{25}{24} \tilde{\mathbf{E}} \cdot \mathbf{t}_{i,m}^p |^n - \frac{1}{12} \tilde{\mathbf{E}} \cdot \mathbf{t}_{i,m}^p |^{n-1} \right. \\ &\quad \left. + \frac{1}{24} \tilde{\mathbf{E}} \cdot \mathbf{t}_{i,m}^p |^{n-2} \right) \Delta_{i,m}^p, \end{aligned} \quad (7.19)$$

$$\begin{aligned} \mathbf{E} \cdot \mathbf{n}_j^d |^{n+1} &= \mathbf{E} \cdot \mathbf{n}_j^d |^n \\ &+ \frac{\Delta t}{\tilde{\epsilon}_j^d A_j^d} \sum_n \left(\frac{25}{24} \tilde{\mathbf{H}} \cdot \mathbf{t}_{j,k}^d |^{n+\frac{1}{2}} - \frac{1}{12} \tilde{\mathbf{H}} \cdot \mathbf{t}_{j,k}^d |^{n-\frac{1}{2}} \right. \\ &\quad \left. + \frac{1}{24} \tilde{\mathbf{H}} \cdot \mathbf{t}_{j,k}^d |^{n-\frac{3}{2}} \right) \Delta_{j,k}^d \\ &+ \Delta t \frac{\tilde{\sigma}_j^d}{2 \tilde{\epsilon}_j^d} (\mathbf{E} \cdot \mathbf{n}_j^d |^{n+1} + \mathbf{E} \cdot \mathbf{n}_j^d |^n), \end{aligned} \quad (7.20)$$

where $\Delta_{i,m}^p$ are the lengths of the primary edges $\mathbf{t}_{i,m}^p$ and $\Delta_{j,k}^d$ are the lengths of the dual edges $\mathbf{t}_{j,k}^d$. Note that the notation $\mathbf{E} \cdot \mathbf{n}_j^d |^n$ means $\mathbf{E} \cdot \mathbf{n}_j^d$ at time step n and accordingly for the other components. The ABS3 method is chosen due to its superior stability properties compared to the normally used leap-frog scheme, see Section 7.4.4.

7.4.3 Preservation of divergence

When solving the Maxwell equations it is common to solve only for Faraday’s and Ampère’s laws. However, the Maxwell equations also include two divergence relations (Gauss’ laws) for the electric flux density vector, \mathbf{D} , and the magnetic flux density vector, \mathbf{B} , that have to be satisfied in order to ensure a physical solution. The divergence of the \mathbf{B} -field is always equal to zero, while the divergence of the \mathbf{D} -field is zero in vacuum. For linear, non-dispersive, isotropic media the simple relationships between the fields are given by

$$\mathbf{B} = \mu \mathbf{H}, \quad \mathbf{D} = \epsilon \mathbf{E}. \quad (7.21)$$

Consider,

$$\begin{aligned} \frac{\partial}{\partial t} \int_{V_q^p} \nabla \cdot \mathbf{B} dV &= \frac{\partial}{\partial t} \int_{V_q^p} \nabla \cdot (\mu \mathbf{H}) dV = \frac{\partial}{\partial t} \oint_{A^p} \mu \mathbf{H} \cdot \mathbf{n}^p dA \\ &= \sum_i \tilde{\mu}_i^p \frac{\partial}{\partial t} \int_{A_i^p} \mathbf{H} \cdot \mathbf{n}_i^p dA = - \sum_i \sum_m \int_{\Gamma_{i,m}^p} \tilde{\mathbf{E}} \cdot \mathbf{t}_{i,m}^p dl = 0, \end{aligned}$$

where Gauss' divergence theorem and (7.13) have been used. The last expression is zero because in the double sum the electric field projected along the edges belonging to the faces i appear twice with opposite signs. Hence, the divergence of the magnetic flux density vector is preserved up to machine precision. The proof for the divergence of the electric flux density vector is analogous. Note that the proof does not assume that the faces have to be planar, which is not necessarily the case for the dual faces.

7.4.4 Stability analysis of the FV solver

An important issue for explicit solvers across many disciplines is how to obtain long-term stability. In 2D we have shown in experiments that our FV solver is stable if ABS3 is used for the time discretization [28]. The stability region for ABS3 in the scalar case is given in the left part of Figure 7.4. The great advantage of using ABS3 compared to the commonly used leap-frog scheme is that we obtain a stable and accurate solution even if the highest frequency eigenvalues of the space operator have a small real part. Using ABS3 will suppress the amplitude of the highest frequency components, but they are not accurately resolved by the grid and are therefore considered parasitic.




Figure 7.4. The stability region for ABS3 is shown to the left. To the right the eigenvalues for the perturbed $4 \times 4 \times 4$ grid. Note the different x-scales in the figures.

To analyze the eigenvalues in vacuum for the semi-discrete case, we write the solver on matrix form as

$$\frac{\partial}{\partial t} \begin{pmatrix} \mathbf{H} \cdot \mathbf{n}^p \\ \mathbf{E} \cdot \mathbf{n}^d \end{pmatrix} = \begin{pmatrix} \mathbf{0} & \mathcal{A} \\ \mathcal{B} & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{H} \cdot \mathbf{n}^p \\ \mathbf{E} \cdot \mathbf{n}^d \end{pmatrix}, \quad (7.22)$$

where the operators \mathcal{A} and \mathcal{B} take care of the space discretization of the respective fields.

Consider a grid consisting of $4 \times 4 \times 4$ cubes, where each cube is split into five tetrahedra in the same way as the cells in the transition layer, see Figure 7.6. Calculating the eigenvalues numerically to the block-matrix above reveals that they are all on the imaginary axis. In this case the primary and dual grids are mutually orthogonal. Let us perturb the grid such that the orthogonality is lost by moving the midpoint along one of the Cartesian directions. The eigenvalues, scaled with the CFL number, are shown in the right part of Figure 7.4. Hence, the loss of orthogonality locally is enough to get eigenvalues with a nonzero real part.

For the fully-discrete case, let

$$\mathbf{z}^n = \left(\mathbf{E} \cdot \mathbf{n}^d |n \quad \mathbf{H} \cdot \mathbf{n}^p |n^{-\frac{1}{2}} \quad \mathbf{E} \cdot \mathbf{n}^d |n^{-1} \quad \mathbf{H} \cdot \mathbf{n}^p |n^{-\frac{3}{2}} \quad \mathbf{E} \cdot \mathbf{n}^d |n^{-2} \right)^T. \quad (7.23)$$

After some straightforward rearrangements we are able to write (7.19) and (7.20) on matrix form as $\mathbf{z}^{n+1} = \mathcal{P}(\mathcal{A}, \mathcal{B})\mathbf{z}^n$, where

$$\mathcal{P} = \begin{pmatrix} \mathcal{I} + \frac{625}{576} \Delta t^2 \mathcal{B} \mathcal{A} & \frac{23}{24} \Delta t \mathcal{B} & -\frac{25}{288} \Delta t^2 \mathcal{B} \mathcal{A} & \frac{1}{24} \Delta t \mathcal{B} & \frac{25}{576} \Delta t^2 \mathcal{B} \mathcal{A} \\ \frac{25}{24} \Delta t \mathcal{A} & \mathcal{I} & -\frac{1}{12} \Delta t \mathcal{A} & \mathbf{0} & \frac{1}{24} \Delta t \mathcal{A} \\ \mathcal{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathcal{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathcal{I} & \mathbf{0} & \mathbf{0} \end{pmatrix}. \quad (7.24)$$

Analyzing the eigenvalues of the companion matrix \mathcal{P} , for the perturbed $4 \times 4 \times 4$ grid, results in the eigenvalue spectrum shown in Figure 7.5, where all eigenvalues are within the unit circle. The time step is chosen on the stability limit for the FV solver and is given by the largest eigenvalue of the semi-discrete problem. For a general unstructured grid an eigenvalue analysis is not possible and hence the CFL condition is not known. A rule of thumb, which works well in practice, is to decrease the time step given by the stability condition for stand-alone FD-TD with a factor given by the quotient between the shortest primary edge and the shortest edge in the structured grid. The eigenvalues close to the origin in Figure 7.5 have to do with the fact that ABS3 is a multi-step scheme, but only one of the roots are relevant and we get a number of irrelevant roots. However, because these roots are all close or equal to zero they are quickly damped away.

If ABS3 is replaced by the leap-frog scheme and the same time step is used, the largest eigenvalue is of the order 1.004, despite the fact that this time step is well within the stability limit for leap-frog. Hence, not surprisingly the leap-frog scheme is unstable due to the fact that we have eigenvalues with a nonzero real part.

However, using ABS3 does not guarantee long-term stability on general tetrahedral grids in 3D. Eigenvalues to the semi-discrete problem might have a real part that is too large for ABS3 to handle. One way of enhancing the numerical stability is to suppress the amplitude of the highest frequency components. A low-pass filter is ideal for this purpose, the undesirable highest frequency components are eliminated, whereas the rest of the solution is remained unchanged. We use a spatial




Figure 7.5. Eigenvalues of the companion matrix \mathcal{P} using ABS3 for the perturbed $4 \times 4 \times 4$ grid. The time step is chosen on the stability limit.

filter based on the Laplace operator. The filter is only applied to the node values and therefore it does not act on edges, which are orthogonal to their corresponding faces. For edges which are not orthogonal the larger the deviation from orthogonality the more the filter affects the edge values. A filtered value $\hat{\mathbf{H}}_r$ of \mathbf{H}_r is given by

$$\hat{\mathbf{H}}_r = \mathbf{H}_r + \frac{1}{2N} \sum_{q=1}^N (\mathbf{H}_q - \mathbf{H}_r), \quad (7.25)$$

where the sum is taken over the nodes q , which share an edge with node r . The filtering for the electric node values is analogous. This filter improved the stability considerably for the grids used in Section 7.7 without losing accuracy in the solution. Note that this filter is not needed in 2D (Chapter 6), where the use of ABS3 is sufficient to obtain long term stability. However, the grids in 3D are in general much more complex for realistic applications and the deviation from orthogonality could be large, which we believe is the main source of the stability problems.

7.5 Hybridization

The 3D hybridization technique requires the unstructured grid to have a transition layer consisting of semi-structured cells. This means that the outermost unstructured cells are obtained by splitting a brick into five tetrahedra (the brick is of the same size as the FD-TD cell). This outermost layer of cells is called the transition layer because the unstructured grid and the structured grid overlaps here, see Figure 7.6. Hence, a number of electric field components coincide in space and time for the solvers in the transition layer and it is these components that are involved in the actual communication between the solvers.

Based on Figure 7.6 the hybridization technique can be described as follows:




Figure 7.6. The transition layer is based on coincidental structured and unstructured cells. To the left an FD-TD cell and to the right a cluster of five tetrahedra forming a brick is shown.

Left part of Figure 7.6,
the structured cell

The four Cartesian components to the left (filled) are sent to the corresponding components in the unstructured region. The diagonal component (also filled) is calculated by interpolation of the four filled components and is also sent to the unstructured region.

↓
(A full time step is taken)

The remaining eight Cartesian components (hollow) are received from the corresponding components in the unstructured region.

Right part of Figure 7.6,
the unstructured cells

The five filled components are received from the structured region and are treated by the unstructured solver as an inhomogeneous Dirichlet boundary condition.

↓
(A full time step is taken)

When the components are updated in the unstructured region the remaining eight Cartesian components (hollow) are sent to the structured solver.

The remaining five diagonal components (ordinary arrows) in the right part of Figure 7.6 are not involved in the hybridization. Note that sending the eight Cartesian components from the unstructured solvers is equivalent in space to sending the four tangential \mathbf{H} -fields in the transition layer to FD-TD.

It is worthwhile to point out that the primary and dual grids are mutually orthogonal in the transition layer. Without this orthogonality also node values would have to be sent from the structured region, to update the electric edge values in the transition layer according to (7.16). In that case a careful interpolation would have to be performed because there are no nodal values in the FD-TD solver. That could possibly jeopardize the accuracy and stability of the hybrid solver.

The transition layer is one of the keys to the success of the hybrid scheme

because it gives a well defined communication between the structured and unstructured solvers. One could think of methods without a transition layer but then the components used for communication are not (in general) collocated and hence interpolation is needed which will deteriorate the hybridization. The one cell overlap also makes it possible to communicate only \mathbf{E} -fields between the solvers and hence fit into a multi block concept better because the FD-TD code is only affected once per time step.

7.6 Stability of the hybrid methods

In this section we discuss the stability of the hybrid solvers. The implicit FE solver is unconditionally stable as a stand alone code. The FD-TD solver is also stable as long as the CFL condition in (4.10) is not violated. The explicit FV-TD solver is stable on orthogonal grids but on general grids it may exhibit weak instabilities, see Section 7.4.4.

For open problems, we demonstrate in Section 7.7 that it is possible to obtain very good results using the hybrid solvers. But for some cavity problems the hybrid solvers suffer from instabilities caused by the communication between the solvers. Cavity problems are often found in electromagnetic compatibility problems.

Two similar cavity problems are solved here to demonstrate the instabilities that can arise. In the first problem a closed vacuum cavity problem is studied. In the second problem a slits is introduced in the cavity. The cavity has a rectangular shape of $16 \times 8 \times 8$ m and an FD-TD grid with $16 \times 8 \times 8$ cells is used (cell size 1 m). A point source in $E_z(5, 5, 5)$ is used for excitation of the cavity. The pulse shape is a derived Gaussian, but as will be discussed later, the pulse shape is not crucial. The simulation is probed in $E_z(5, 11, 5)$. In all cases presented the CFL number is 0.433 which means that it takes four time steps to propagate the solution one meter.

The basic problem setup is displayed in Figure 7.7. For the test of the hybridization technique half of the cavity volume is replaced by a semi-structured tetrahedral grid in which the FD-TD cells are split into five tetrahedra (as in the right part of Figure 7.6). The hybrid interface divides the cavity in two blocks of cells.

In the first problem there is no physical damping because the volume is lossless and hence the cavity is infinitely resonant. This is also seen in the FD-TD response where the amplitude in the probed point is essentially constant, see Figure 7.8. In contrast to the pure FD-TD simulations the hybrid solutions become unstable. Four different hybrid results are obtained, see Figure 7.8. First the FD-FV hybrid is used which becomes unstable after roughly 30 000 time steps. The same happens if we use the FD-FE hybrid utilizing the Θ -method with $\Theta = 0.5$ (see Section 6.3.3). In the third case the FD-FE hybrid is used but with $\Theta = 0.51$. This gives a slightly dissipative timestepping scheme but still the solution becomes unstable after approximately 60 000 time steps. The FD-FE hybrid is also used in the fourth case but this time the BDF2 timestepping scheme is used (see Section 6.3.3). In




Figure 7.7. $16 \times 8 \times 8$ FD-TD cells constitute the basic cavity problem. The outer walls are perfectly electric conducting and excitation and probing is done in $E_z(5, 5, 5)$ and $E_z(5, 11, 5)$. The hybrid interface divides the cavity into the left part where FD-TD is used and the right part where the $8 \times 8 \times 8$ cells are split into tetrahedra to constitute the unstructured region. In the second cavity problem a slits is introduced (illustrated by the shadowed rectangle).

this case the solution does not becomes unstable. The trend shown in Figure 7.8 continues up to 900 000 time steps where a plateau level of 10^{-7} is reached. After 1 500 000 time steps the amplitude is still 10^{-7} .




Figure 7.8. The infinitely resonant cavity (left) exhibits unstable hybrid solutions for the FD-FV hybrid as well as for the three FD-FE hybrid methods. The difference between the three FD-FE hybrid methods is the choice of timestepping scheme; the Θ -method with $\Theta = 0.5$ (Crank-Nicholson), $\Theta = 0.51$, and the BDF2-method. The BDF2 solution does not exhibit instabilities up to 700 000 time steps. To the right the response for the slitted cavity is shown. In this case the unstable solutions for all four hybrid methods are almost the same.

The FE solver is slightly dissipative if $\Theta > 0.5$ or the BDF2 method is used. This gives rise to energy losses which is clearly seen in Figure 7.8. These timestepping errors are decreased if we pay the price of taking a smaller timestep. If we lower the CFL number by a factor of two the solution for $\Theta = 0.51$ has a one decade amplitude loss over 150 000 time steps whereas the BDF2 solution has a one decade

amplitude loss over 600 000 time steps. (Remember that these two numbers are obtained with a time step that is half of that used in Figure 7.8.) The $\Theta = 0.51$ solution exhibit a weaker instability if the CFL number is halved. But the BDF2 solution shows a stronger instability if the CFL number is halved.

The strength of the instabilities are different for the different hybrid solvers and it is interesting to follow the instability curves backwards in time to $t = 0$. Even though the strength of the instabilities are different it is clear that they all originate from the round off error level of 10^{-16} which is the dynamic range for double precision floating point numbers in this simulation. This means that even if the excitation itself does not contain the dominant unstable mode the round off errors will inevitable trig the unstable modes.

In the second test case the cavity is opened which introduces a physical loss mechanism. The aim of the test is to show that also highly but not infinitely resonant cavities might become unstable. A slits of two cells width is opened on top of the yz -plane at the lower x -end of the cavity. The slits is illustrated by the shadowed rectangle in Figure 7.7. Outside the cavity a surrounding layer of FD-TD cells are used with a simple absorbing boundary condition. The energy loss through the slits is characterized by a one decade amplitude decrease over 200 000 time steps

This slitted cavity problem also exhibit instabilities when the hybrid methods are used. The instabilities appear after only 6 000 time steps and all four hybrid methods have approximately the same response to the dominant unstable mode. If we follow the instability curves back to $t = 0$ we see that the unstable mode is triggered by the excitation on a level of 10^{-5} . In this case the dominant unstable mode is almost unaffected by the choice of timestepping methods used for the FE solver. If the CFL number is halved the instability is somewhat weaker but still affects the four hybrid schemes approximately the same.

Instabilities for FD-FE hybrids have been observed and addressed by other researchers [51]. Our efforts include various spatial low-pass filters in FD-TD close to the hybrid interface, different dissipative timestepping schemes, and splitting the transition layer cubes into six tetrahedra instead of five. None of these efforts have been fully successful to guarantee general stability. However, in [87, 88] Rylander et al. present a stable FD-FE method. No rigorous stability proof is given but the argumentation is based on properties for symmetric matrices. Pyramidal elements and cube elements are used to join the unstructured tetrahedral grid with the structured hexahedral FD-TD grid and thus the need for interpolating diagonal components is avoided.

7.7 Results

In this section the convergence of the hybrid code is studied numerically for a few simple vacuum cases. The hybrid code is also validated in four different scattering cases, a PEC sphere, a dielectric sphere, the NASA almond and the generic aircraft *RUND*. For the PEC and dielectric spheres analytical solutions (Mie series) for the radar cross section (RCS) are known in frequency domain. For the almond and the generic aircraft, converged Method of Moments solutions are used as reference solutions. Since the computation of RCS requires the scattered field intensity at infinity, we use a near-to-far-field transformation [63] to compute the far-field solution. The computational domains are truncated by eight layers of U-PML cells. The CFL number is set to 0.2 in all cases, where CFL equal to one is the stability limit for FD-TD, see (4.11). The grids used for stand alone FD-TD are generated such that all edges which belong to a cell whose centroid is inside the object are considered to be PEC or dielectric edges. The center of the sphere coincides with the center of an FD-TD cell.

7.7.1 Convergence in vacuum

Procedure

We estimate the order of convergence for vacuum test cases. Since we have analytical solutions for these cases, we only need two numerical solutions: one coarse grid solution and one fine grid solution. Since we consider the electric components in the grids to be located at the center of each edge, and we want to avoid spatial interpolation, the fine grid must have edges that are three times shorter than the edges in the coarse grid. Furthermore, we use an odd number of cells in the z-direction and an even number of cells in the x- and y-directions. This will ensure that we have a z-edge in the center of the unstructured grids.

In the structured grid we use cubic cells with side length $\Delta = 1$ m for the coarse grid and $\Delta = \frac{1}{3}$ m for the fine grid. The unstructured grids have a size of $4 \text{ m} \times 4 \text{ m} \times 3 \text{ m}$ and they are inserted in the center of an FD-TD grid that is $14 \text{ m} \times 14 \text{ m} \times 13 \text{ m}$.

We probe one or more electric field components every time step on the coarse grid. On the fine grid we probe every third time step. We always probe electric field components because their values are represented at $t = n\Delta t$ where $n = 0, \dots, N_t$ and hence we need no temporal interpolation.

We calculate the convergence estimate p_{max} as

$$p_{max} = \frac{\ln(\max(\text{abs}(err_c))) - \ln(\max(\text{abs}(err_f)))}{\ln(3)}. \quad (7.26)$$

where err_c and err_f are time vectors containing the errors in a probed point. The indexes c and f stands for coarse and fine. The convergence estimate p_{mean} is

calculated equivalently as

$$p_{mean} = \frac{\ln(\text{mean}(\text{abs}(\text{err}_c))) - \ln(\text{mean}(\text{abs}(\text{err}_f)))}{\ln(3)}. \quad (7.27)$$

Stand alone unstructured solvers

First we study the convergence properties of the unstructured solvers in stand alone mode. This is done using semi-structured unstructured grids, where each cube is split into five tetrahedra (as in the right part of Figure 7.6). We excite the unstructured grid by supplying analytical boundary conditions, i.e. a time dependent Dirichlet boundary condition. This is done in the same way as the FD-TD solver sends values to the unstructured solvers. Hence, it is only the the four (filled) Cartesian components in the left part of Figure 7.6 that are given boundary values. The filled diagonal component is calculated by interpolating the four Cartesian components.

For the incident field we use a Gaussian shaped pulse propagating in x-direction with electric polarization in z-direction. The half-maximum-pulse-width corresponds to $11\Delta_c$.

For both methods we use $CFL = \sqrt{3}/4$ where the CFL number is defined as in (4.11). We use the Crank-Nicholson method for the FE timestepping.

Our first probe is located at the centermost z-edge in the unstructured grid. The second probe is located with an offset of $(1, -1, 0)$ m to the first probe.

Table 7.1 contains the convergence estimates for the two probes. We have almost exactly two in all cases and may hence conclude that the unstructured solvers have second-order accuracy on semi-structured grids.

Method	FE		FV	
Edge	1	2	1	2
p_{mean}	2.00	2.00	1.99	1.99
p_{max}	1.99	1.99	1.99	1.99

Table 7.1. Estimates of the order of convergence for the two unstructured grid methods. Edge 1 is the centermost edge.

Hybrid solvers

We will now study the hybrid code. We use the same two semi-structured grids and the same incident field as in the previous test case. The incident field is generated by Huygens' surfaces which are located two meters from the outer boundary. For FD-TD and the FD-FE hybrid we use $CFL = \sqrt{3}/2$ and for the FD-FV hybrid we use $CFL = \sqrt{3}/4$. A highly absorbing ABC is used (a twelve cells thick PML layer) to ensure that reflections from the outer boundary do not noticeably affect our convergence estimates.

The absolute value of the errors for E_z at $(x, y, z) = (4, 7, 6.5)$ are shown in Figure 7.9. We see that they are significantly larger for the hybrid methods.




Figure 7.9. Errors for two refinement levels and three methods for E_z at $(x, y, z) = (4, 7, 6.5)$ (thick lines: coarse grid, thin lines: fine grid).

We probe the solution in all E_z components (on the coarse grid) that lie one meter or one and a half meter from the outer boundary of the transition layer. We calculate the l_2 error and use these time dependent vectors to estimate the convergence. Estimates are given in Table 7.2. Clearly, we have super-linear convergence. However, we do not have second-order convergence.

FD-TD		FD-FE		FD-FV	
p_{mean}	p_{max}	p_{mean}	p_{max}	p_{mean}	p_{max}
1.96	1.98	1.25	1.11	1.15	1.04

Table 7.2. Estimates of the order of convergence for the two hybrid methods and the FD-TD method.

The convergence estimates for the two hybrid methods are rather similar. This fact and the fact that our previous test case shows that the FV-TD and FE-TD solvers are second-order accurate on these grids in stand alone mode, strongly indicates that our hybridization technique destroys the second-order accuracy. The likely culprit is the interpolation of diagonal values. The error in the interpolation itself is proportional to Δ^2 . However the interpolated value is then used in what is basically a difference approximation. If this happened in only one point, it would not destroy the global accuracy. However, because it happens in every diagonal on the outer surface of the transition region, it does affect the global accuracy.

To validate our hypothesis that the interpolation of diagonal values are causing the deteriorated accuracy, we perform two tests. In the first test we replace the

diagonal values with analytical values. The absolute value of the errors for E_z at $(x, y, z) = (4, 7, 6.5)$ are shown in Figure 7.10. Convergence estimates are given in Table 7.3.




Figure 7.10. Errors for two refinement levels for the two hybrid methods for E_z at $(x, y, z) = (4, 7, 6.5)$ when diagonal values are substituted by analytical values. (thick lines: coarse grid, thin lines: fine grid).

FD-FE		FD-FV	
p_{mean}	p_{max}	p_{mean}	p_{max}
1.93	1.96	2.11	1.88

Table 7.3. Estimates of the order of convergence for the two hybrid methods when diagonal values are substituted by analytical values.

In the second test we use higher order interpolation. Four values are used instead of two when performing the interpolation. This procedure is unstable but it seems to be second-order accurate for the brief time span that is unaffected by the instability.

These tests clearly demonstrate that it is the interpolation of diagonal values that destroys the second-order accuracy.

7.7.2 Scattering from a PEC sphere

A standard test case to validate electromagnetic codes in 3D is scattering from PEC spheres. In this case the sphere has a radius of three meters and a Gaussian pulse with horizontal polarization impinges the sphere at an angle of $\theta = 0$, using spherical polar coordinates. A few layers of unstructured cells are used in the vicinity of the sphere to obtain a good geometrical representation of the sphere. A structured grid is used for the rest of the domain. The edges in the unstructured

region are approximately of the same length as the ones in the structured region. The results for monostatic RCS and bistatic RCS at 27 MHz are compared to the analytical Mie series solutions and with three staircased FD-TD solutions in Figure 7.11.




Figure 7.11. Monostatic RCS (left) and bistatic RCS at 27 MHz (right) for a PEC sphere of radius three meters. $\Delta = 1\text{m}$ corresponds to a resolution of eleven points per wavelength at 27 MHz. $\theta = 0$ is the monostatic angle. The FD-FV hybrid solver is used but the FD-FE hybrid solver gives similar results.

The results using the hybrid solver are clearly much better than the ones obtained using staircased FD-TD and agree very well with the analytical solution even with a moderate resolution of eleven points per wavelength. It is interesting to note the behavior of FD-TD at the monostatic angle, $\theta = 0$, in the right part of Figure 7.11. It takes a resolution of 44 points per wavelength to even get close to the correct behavior of the solution. Note that increasing the resolution with a factor of four requires 64 times more memory and 256 times more arithmetic operations in FD-TD, and we are still not able to obtain an accurate solution. Hence, staircasing introduces large geometrical errors and an extremely fine resolution is needed to rectify these errors.

To estimate the order of accuracy for the hybrid solver in this case we perform a grid refinement study. We refine the grids with a factor of three, because

that implies that the staggered components on the coarse grid are located at the same positions on the finer grids. The three grids have edge lengths $\Delta = 0.99$ m, $\Delta = 0.33$ m, and $\Delta = 0.11$ m, respectively. A Gaussian pulse is used as incoming pulse and the pulse width is chosen such that the significant spectral components are well resolved with twenty points per wavelength for the coarse grid. The scattered electric field is sampled at different points in the near zone located both in the backscattering region and in the shadow region behind the sphere. To analyze the order of accuracy in a standard numerical manner we study the difference between the scattered fields on the coarse- and intermediate-grid and on the intermediate- and fine-grid for each sample point. Averaging these differences over all sample points results in the two curves in Figure 7.12. Finally, from these two curves we obtain estimates of the order of accuracy using L_1 norm. The result is given in Table 7.4.

FD-FV	FD-FE
1.85	1.92

Table 7.4. Convergence estimates.

Choosing different sample points and a different norm does not affect the accuracy figures significantly. For a more careful study of the order of accuracy we probably need one more level of refinement.




Figure 7.12. The mean absolute differences in the scattered field in the sample points as a function of time for the FD-FV hybrid solver and the FD-FE hybrid solver.

7.7.3 Scattering from a dielectric sphere

The ability to treat different materials is an important feature of a general purpose electromagnetic code. In this thesis we only consider materials whose properties are independent of frequency. Results for frequency dependent materials are given in [29]. The sphere has the same dimensions as in the previous case and the same incoming Gaussian pulse is used. The relative permittivity, ϵ_r , of the sphere equals four. Tetrahedra are used throughout the inside of the sphere. For a larger sphere tetrahedra should be used only close to the surface of the sphere, and structured cells should be used inside the sphere to reduce the memory requirements and computational time. However, here the sphere is small in terms of wavelengths and the extra computational cost for using tetrahedra throughout the sphere is small.

In Figure 7.13 the results from our hybrid solver are compared to the Mie series and to FD-TD solutions with different resolutions. The frequency is 15 MHz, which implies a resolution of ten points per wavelength in the sphere when $\Delta = 1$ m. This resolution is enough for the hybrid solver to obtain a good correspondence with the Mie series for all angles. However, FD-TD needs a resolution of 40 points per wavelength to get results of similar accuracy.




Figure 7.13. Bistatic RCS for a dielectric sphere, $\epsilon_r = 4$, of radius three meters at 15 MHz (left). $\Delta = 1$ m corresponds to a resolution of ten points per wavelength in the sphere. A zoom of the plot is shown to the right where an FD-TD solution using $\Delta = 1/4$ m is added. The FD-FV hybrid solver is used. Similar results are obtained with the FD-FE hybrid.

7.7.4 Scattering from the NASA Almond

Another common test problem for electromagnetic methods is the NASA almond. It is defined in for instance [24]. Figure 7.14 shows the surface grid on the NASA Almond.




Figure 7.14. The NASA Almond.

The computations are performed at Ericsson Saab Avionics using the FD-FV hybrid. The monostatic RCS for nineteen different angles of incidence are shown in Figure 7.15. They agree very well with a converged MoM solution.

In order to see how large the reflection from the grid interface is, the interior of the NASA almond is filled with tetrahedra and a vacuum computation is performed. In Figure 7.15 these results are labeled FV-TD void. The errors are -60 to -80 dB. This can be compared with the errors for pure FD-TD which are -110 to -125 dB.

The frequency-domain near-to-far-field transform with dispersion compensation is used to generate the far-field. The compensation is based on the dispersion relation for FD-TD. Hence the higher error for the FD-FV hybrid might not only be due to reflections in the transition layer, but could also be partially caused by the discrepancy between the FD-TD dispersion relation and the actual dispersion for the FV method. In order to investigate the the dispersion compensation, we perform computations with $\phi = 180$ with and without dispersion compensation. The results are given in Table 7.5.

Dispersion compensation	Yes	No
FV-TD PEC	-39.2	-39.0
FV-TD void	-67.7	-63.9
FD-TD void	-122.0	-74.9

Table 7.5. Monostatic RCS (σ_θ) for $\phi = 180$.

The PEC RCS results are almost unchanged, while there is a drastic increase in the FD-TD void RCS and a small increase in the FV-TD void RCS. It is clear that the dispersion compensation plays a major roll in the excellent FD-TD void result.




Figure 7.15. Monostatic RCS for the NASA Almond at 3 GHz.

We can also conclude that our hybridization method does not introduce reflections that destroys the accuracy of the monostatic RCS for the NASA almond.

7.7.5 Scattering from the generic aircraft *RUND*

The RCS for canonical objects like spheres are often used for validation because analytical solutions are known. However, from an industrial perspective spheres are not targets of practical interest. A more realistic case is scattering from the generic aircraft *RUND*. This aircraft is approximately one meter long, one meter between the wing tips and half a meter high, see Figure 7.16. The computational domain is discretized using roughly 150 000 tetrahedra and 500 000 FD-TD cells. The FD-TD cells have one centimeter long edges, which corresponds to a resolution of 20 points per wavelength at 1.5 GHz. A Gaussian pulse illuminates the aircraft head on with either vertical polarization or horizontal polarisation, see Figures 7.17 and 7.18.




Figure 7.16. Part of unstructured mesh around generic aircraft *RUND*. The structured mesh continues outside the shown unstructured cells. Note that the interface between the grids is staircased to minimize the number of unstructured cells.




Figure 7.17. Comparison of bistatic RCS between the hybrid solver, FD-TD and a MoM solver at 1.5 GHz for vertical polarization. The resolution with $\Delta = 10$ mm corresponds to 20 points per wavelength. BDF2 is used for timestepping in the FE hybrid.




Figure 7.18. Comparison of bistatic RCS between the hybrid solver, FD-TD and a MoM solver at 1.5 GHz for horizontal polarization. The resolution with $\Delta = 10$ mm corresponds to 20 points per wavelength. BDF2 is used for timestepping in the FE hybrid.

The agreement with the Method of Moments (MoM) solution, which is considered to be a good reference solution, is excellent for all angles except the first 30 degrees for the horizontal polarization, see Figures 7.17 and 7.18. The FD-TD solution with four millimeter long edges is good for some angles, but to get a good agreement for the monostatic angle, $\phi = 0$, where the staircasing effect in this case is at worst, would require a very fine resolution.

The execution times and memory requirements are in this case comparable for the FV hybrid solution, the FE hybrid solution, and the FD-TD solution with four millimeter long edges. However, none of the hybrid solvers are fully optimized in terms of memory requirements and execution times. The FD-FV hybrid would benefit from using local timestepping in the unstructured region and the convergence properties of the FE iterative solver could be improved by a more sophisticated preconditioner. A detailed comparison of the efficiency and accuracy of the methods will be performed after these issues have been rectified.

7.8 Conclusions

We have introduced an explicit as well as an implicit hybrid solver for the Maxwell equations in time domain. By combining the efficiency of the classical FD-TD method with the flexibility of an solver for unstructured grids we obtain a very efficient and accurate solver. Due to the fact that the explicit FV solver must obey a stability limit where the maximum time step is related to the smallest unstructured cell, the explicit solver is intended for problems where the cell sizes in the unstructured region are of the same size as the FD-TD cells. The implicit FE solver has no stability limited regarding the time step and hence it is suitable for grids used to resolve fine details.

The FV solver uses ABS3 for time discretization. Together with a spatial Laplace filter we can run all the FD-FV scattering cases in Section 7.7 without any signs of instability. The Θ -method used as time integrator for the FE solver gives us a possibility to stabilize the FD-FE hybrid. This is not needed in any of the convergence studies, where FD-FE is used. But for the calculation of RCS for the generic aircraft *RUND* stabilizing of the FD-FE hybrid is necessary. Stable results are obtained for $\Theta = 0.54$ and also if BDF2 is used for timestepping.

The numerical examples presented in this paper clearly demonstrate that our hybrid solvers are much more efficient than standard FD-TD on complex geometries. It is also interesting to note that the numerical results show no signs of large reflections at the grid interface. This is clearly illustrated in the Almond test case where the reflections from the hybrid interface is approximately 20 dBsm lower than the scattering from the object itself. We believe that the low level of reflection is mainly due to the transition layer, which ensures a smooth transition from the structured to the unstructured grid. The number of unstructured cells is minimized by allowing a staircased interface between the structured and unstructured grids.

Chapter 8

Thin-Wire Subcell Models

8.1 Introduction to thin-wire subcell models

This chapter deals with the problem of how to include thin wires into the FD-TD method. Thin wires are often important parts in electromagnetic compatibility problems and antenna problems and have been studied since the famous experiments conducted by Hertz [46] more than a century ago. Modern antenna techniques, and a nice overview of the antenna development, is found in Balanis [7].

The most straightforward technique to include wires into the FD-TD method would of course be to model the wire as a perfect electric conductor, and hence set all tangential electric components on the wire surface equal to zero. But this strategy requires a very fine discretization in order to resolve the cross-section of the wire which is clearly out of question in most practical cases. If we use the hybrid methods described in Section 7 we could wrap a layer of unstructured cells around the wire and resolve the cross-section of the wire locally. But this would give us many very small cells and hence the computational burden becomes too heavy. Therefore, thin-wire subcell models have been developed with the purpose of capturing the effects of a wire present in the discretized 3D space without resolving the wire radially.

In this chapter a thin-wire model for FD-TD is extended to allow for an arbitrary orientation of the wire. This is important in order to treat wires that cannot be aligned with the Cartesian grid, for example circular loop antennas and tilted straight wires. The only way to incorporate arbitrarily oriented wires into FD-TD today is to accept a staircased discretization of the wire. However, this leads to large errors, and the wire length does not converge if a finer resolution is used. This implies that resonance frequencies, and related quantities, can be erroneous with tens of percent.

There are basically two approaches to construct thin-wire subcell models for FD-TD. Holland et al. [49] separates the wire from the 3D volume and solves two

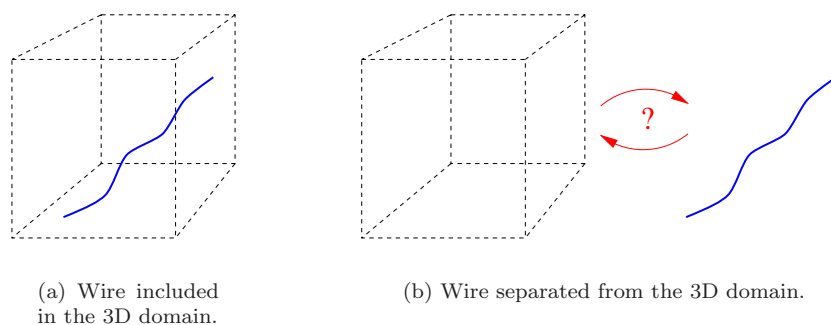


Figure 8.1. Separation of the two problems.

mutually interacting problems, see Figure 8.1. The idea behind this wire model is based on solving partial differential equations governing the current and charge density on the wire. Umashankar et al. [103] does not solve a wire problem explicitly but includes the effect of the wire directly into the stencil of the Yee scheme. This is done by using the contour-path interpretation of the Maxwell-Faraday law [98].

There are several other techniques to solve problems including thin wires than to use subcell models in FD-TD. The most well known method is the frequency domain electric field integral equation (EFIE). This formulation is often found in the wide spread Method of Moments (MoM) codes such as NEC [75]. Time domain integral equations are also used to formulate a marching-on-in-time procedure to solve wire problems, see for example [12]. The Transmission-line modeling technique (TLM), which is similar to FD-TD, has also been extended to include thin wires, see [86]. For unstructured mesh time domain solvers, Riley et al. have developed thin wire models for both Finite Volumes [85] and Finite elements [83]. There are also hybrid techniques presented combining the MoM in time domain and FD-TD, see for example [14].

Several extensions to the basic wire models for volumetric solvers are presented in the literature. Based on the FD-TD subcell model of Holland et al. [49], Bérenger developed a multiwire formalism to include bundled wires [10]. A multiconductor model is also presented by Włodarczyk et al. for TLM [107]. These models for several wires running close together treat a very common situation in electromagnetic compatibility problems. The method presented by Umashankar et al. [103] is also enhanced by several groups. Douglas et al. [26] have improved the model to enable better treatment of end effects. The basic assumption of a $\frac{1}{r}$ dependency of the normal electric field and tangential magnetic field, along the wire, is improved by adjusting the dependencies at the wire ends. In [106] Watanabe et al. presents an improved feeding gap model giving more accurate impedances compared to the regular “one-cell gap”. Bingle et al. [11] have extended the wire model of Umashankar et al. to also incorporate non-perfect conducting wires.

As will be shown later, the subcell model of Holland et al. [49] is suitable to start with when constructing a wire model that can handle arbitrarily oriented wires.

The subcell model of Umashankar et al. [103] will not be further covered here even though it is commonly used. The formulation requires the wire to be aligned to the Cartesian grid and it is not clear how to generalize this model to arbitrarily oriented wires.

None of the two approaches for FD-TD can be completely defined in continuous space and then discretized. The coupling of the wire and the surrounding field depends on the discrete cell size in relation to the wire radius. This is somewhat bothering but is not yet successfully circumvented. A consequence of this is that a classical order of accuracy analysis is hard to perform. The parameters in the models have basically a logarithmic dependency on the ratio between the cell size and the wire radius and hence, changing the cell size will alter the coefficients in the model. Moreover, if a grid is sufficiently refined, the wire cannot be regarded thin and the model breaks down. For problems where the wire radius is moderate the hybrid techniques described in Chapter 7 might successfully be used to resolve the wire as a perfect conducting object.

This chapter consists primarily of five parts. In Sections 8.2, 8.3, and 8.5 the wire equations are established and discretized. These sections contain essentially known material, and follow [49] quite closely. Sections 8.4 and 8.6 include an extended interpretation to the coupling between the wire equations and the Maxwell equations. The interpretation is new all though the underlying equations are well known. Section 8.7 describes the new shell interpolated tri-linear interpolation technique which can handle arbitrarily oriented thin wires. To our knowledge, there is no open literature that describes this capability for an FD-TD thin-wire model. A stability analysis for a periodic model problem is provided in Section 8.8 where the thin-wire model is shown to be stable under a time step limit. In Section 8.10 results are presented showing that the new shell interpolated tri-linear interpolation technique gives very good results for dipole antennas and circular loop antennas. Finally, some comments on the physical assumptions as well as conclusions are given in Sections 8.9 and 8.11

8.2 Physics

In order to split the original field-and-wire problem we derive a suitable wire equation by studying an infinite long cylinder running in the z-direction, see Figure 8.2.

Following Holland et al. [49] we start by taking the θ component of Faraday's law written in cylindrical coordinates. Together with the boundary condition on the wire surface (tangential component of the electric field is equal to zero on a




Figure 8.2. Cylindrical coordinates.

metallic surface) we have

$$\begin{aligned} \frac{\partial E_r}{\partial z} - \frac{\partial E_z}{\partial r} &= -\mu \frac{\partial H_\theta}{\partial t}, \\ E_z &= 0, \quad r = a. \end{aligned} \quad (8.1)$$

Integration of (8.1) from $r = a$ to R gives

$$\begin{aligned} \frac{\partial}{\partial z} \int_a^R E_r dr - \underbrace{\frac{\partial}{\partial r} \int_a^R E_z dr}_{E_z(R) - E_z(a)} &= -\mu \frac{\partial}{\partial t} \int_a^R H_\theta dr \\ \Leftrightarrow \\ E_z(R) &= \mu \frac{\partial}{\partial t} \int_a^R H_\theta dr + \frac{\partial}{\partial z} \int_a^R E_r dr. \end{aligned} \quad (8.2)$$

The r component of Ampère's law written in cylindrical coordinates with periodic boundary condition for the H_z component reads

$$\begin{aligned} \frac{1}{r} \frac{\partial H_z}{\partial \theta} - \frac{\partial H_\theta}{\partial z} &= \epsilon \frac{\partial E_r}{\partial t}, \\ H_z(r, 0) &= H_z(r, 2\pi). \end{aligned} \quad (8.3)$$

Integration of (8.3) from $\theta = 0$ to 2π gives

$$\begin{aligned} \frac{1}{r} \underbrace{\frac{\partial}{\partial \theta} \int_0^{2\pi} H_z d\theta}_{H_z(2\pi) - H_z(0) = 0} - \frac{\partial}{\partial z} \int_0^{2\pi} H_\theta d\theta &= \epsilon \frac{\partial}{\partial t} \int_0^{2\pi} E_r d\theta \\ \Leftrightarrow \\ -\frac{\partial}{\partial z} \int_0^{2\pi} H_\theta d\theta &= \epsilon \frac{\partial}{\partial t} \int_0^{2\pi} E_r d\theta. \end{aligned} \quad (8.4)$$

In order to close the system of equations formed in (8.2) and (8.4) we have to make some assumptions. These assumptions are made in the borderland between the discretized problem and the (real) continuous problem which is somewhat annoying.

8.2.1 First assumption

We need relations between the field quantities and the wire quantities and hence we make the following quasi-static assumption. Assume that the interaction between the surrounding field and wire have a much smaller time variation than space variation. In that case we can use Biot-Savart's law (8.5) and Coulombs law (8.6) to obtain relations between the field-variables and the wire-variables.

Biot-Savart's laws for direct current (frequency = 0) in an infinite long wire is,

$$H_\theta(r) = \frac{I}{2\pi r}, \quad (8.5)$$

and Coulombs law for static charges on an infinite long wire is,

$$E_r(r) = \frac{Q}{2\pi r\epsilon}, \quad (8.6)$$

where I is the wire-current and Q is the charge per unit length on the wire surface.




Figure 8.3. Wire-variables.

Now we can substitute (8.5) and (8.6) into (8.4),

$$-\frac{1}{2\pi r} \frac{\partial I}{\partial z} \int_0^{2\pi} d\theta = \frac{\epsilon}{2\pi r\epsilon} \frac{\partial Q}{\partial t} \int_0^{2\pi} d\theta \quad \Leftrightarrow \quad \frac{\partial Q}{\partial t} + \frac{\partial I}{\partial z} = 0 \quad (8.7)$$

and (8.5) and (8.6) into (8.2),

$$\begin{aligned} E_z(R) &= \frac{\mu}{2\pi} \frac{\partial I}{\partial t} \int_a^R \frac{1}{r} dr + \frac{1}{2\pi\epsilon} \frac{\partial Q}{\partial z} \int_a^R \frac{1}{r} dr \\ &\Leftrightarrow \\ E_z(R) &= \frac{\mu}{2\pi} \ln\left(\frac{R}{a}\right) \frac{\partial I}{\partial t} + \frac{1}{2\pi\epsilon} \frac{\partial Q}{\partial z} \ln\left(\frac{R}{a}\right). \end{aligned} \quad (8.8)$$

Reformulating (8.8) gives

$$E_z(R) = L \left(\frac{\partial I}{\partial t} + \frac{1}{\epsilon\mu} \frac{\partial Q}{\partial z} \right), \quad R > a, \quad (8.9)$$

$$E_z(R) = 0, \quad R \leq a, \quad (8.10)$$

where L is the inductance per unit length of the wire in Henry/meter $[\frac{Vs}{Am}]$,

$$L(R) = \frac{\mu}{2\pi} \ln\left(\frac{R}{a}\right). \quad (8.11)$$

8.2.2 Second assumption

Equation (8.9) is not a constant coefficient PDE because $L = L(R)$. Since L dictates the influence of the external electric field on the current we must be careful when discretizing (8.8). If for example a discrete current node and a tangential electric component are collocated the inductance per unit length is infinite if calculated using (8.11).

We would like to avoid variable coefficients and hence we make our second assumption to obtain a constant coefficient problem. The assumption will lead us to formulas of the type found in (8.16).

By taking the discrete cell size into account we can determine L . This can be done in several ways and the most obvious choice is to pick a radius, for example $R = \Delta/2$ if the FD-TD cells are cubic ($\Delta_x = \Delta_y = \Delta_z = \Delta$). An argument for this choice is that a discrete field component also represents the field averaged over a surface of $\Delta \times \Delta$. This yields

$$L_{\odot} = \frac{\mu}{2\pi} \ln\left(\frac{\Delta}{2a}\right). \quad (8.12)$$

Another choice is to take an average of (8.11) over some area, see left part of Figure 8.4. Holland et al. [49] presents two strategies where the simplest one for




Figure 8.4. Averaging $\log\left(\frac{R}{a}\right)$ over an equivalent area.

cubic FD-TD cells is to average over a circle with the same area as four times the

FD-TD face area, see right part of Figure 8.4. This gives

$$\begin{aligned} L_o &= \frac{\int_0^{2\pi} \int_a^R r \frac{\mu}{2\pi} \ln\left(\frac{r}{a}\right) dr d\theta}{\int_0^{2\pi} \int_0^R r dr d\theta} = \frac{\int_0^{2\pi} d\theta \frac{\mu}{2\pi} \int_a^R r \ln\left(\frac{r}{a}\right) dr}{\int_0^{2\pi} d\theta \int_0^R r dr} \\ &= \frac{\mu}{2\pi} \left(\ln\left(\frac{R}{a}\right) - \frac{1}{2} + \frac{a^2}{2R^2} \right) \end{aligned} \quad (8.13)$$

with $R = \sqrt{\frac{4}{\pi}} \Delta$ we get

$$L_o = \frac{\mu}{2\pi} \left(\ln\left(\frac{2\Delta}{\sqrt{\pi}a}\right) - \frac{1}{2} + \frac{a^2\pi}{8\Delta^2} \right).$$

We could think of integrating from 0 to R in the radial integral in the numerator of (8.13) but since $E_z = 0$ for $r < a$, see (8.10), (8.13) is justified. Furthermore, we can argue about using the same radial integration limits in the denominator, but Holland et al. [49] reports on better agreements to known analytic solutions to simple problems by using $r = 0$.

A more refined calculation of L , which also holds for non-cubic cells, is found by integrating (8.11) over the four rectangles in Figure 8.5 and divide by the area. We have,




Figure 8.5. Arbitrary rectilinear cells.

$$L_{\square} = \frac{\mu}{2\pi} \frac{\int \int_{\Omega-W} \ln\left(\frac{r}{a}\right) dx dy}{\int \int_{\Omega} dx dy}, \quad (8.14)$$

where Ω is the domain of the four rectangles and W is the wire cross-section. W is removed in the numerator to avoid a logarithmic singularity and is justified by (8.10). As in (8.13) the integral in the denominator is taken over the whole area. Hence, it equals $4AB$. Working through the evaluation of the numerator integrals in (8.14) gives the averaged inductance per unit length,

$$L_{\square} = \frac{\mu}{4\pi} \left(\ln\left(\frac{A^2 + B^2}{a^2}\right) - 3 + \frac{A}{B} \arctan\left(\frac{B}{A}\right) + \frac{B}{A} \arctan\left(\frac{A}{B}\right) \right) + \frac{a^2}{16AB}. \quad (8.15)$$

In the case of quadratic cell cross-sections we can compare L_{\square} , L_{\circ} and L_{\circledast} :

$$\begin{aligned} L_{\square} &= \frac{\mu}{2\pi} \left(\ln\left(\frac{\sqrt{2}\Delta}{a}\right) - \frac{3}{2} + \frac{\pi}{4} + \frac{\pi a^2}{8\Delta^2} \right), \\ L_{\circ} &= \frac{\mu}{2\pi} \left(\ln\left(\frac{2\Delta}{\sqrt{\pi}a}\right) - \frac{1}{2} + \frac{\pi a^2}{8\Delta^2} \right), \\ L_{\circledast} &= \frac{\mu}{2\pi} \ln\left(\frac{\Delta}{2a}\right). \end{aligned} \quad (8.16)$$

From Figure 8.6 there is apparently not much of a difference between L_{\square} and L_{\circ} for quadratic cross-sections. The normalized difference is only ≈ 0.011 for all $\frac{a}{\Delta}$. It




Figure 8.6. Normalized inductance per unit length ($\frac{2\pi}{\mu}L$) for quadratic cross-sections for L_{\square} , L_{\circ} and L_{\circledast} . In this Figure L_{\square} and L_{\circ} are indistinguishable.

is only for larger wire radii, say $a = 0.1\Delta$, that the difference exceeds one percent. We also see that $L_{\circledast} \leq 0$ for $a \geq \frac{1}{2}\Delta$, but L_{\square} is always positive and $L_{\circ} = 0$ for $a \approx 1.13\Delta$. However, it is doubtful whether you can call a wire with a radius equal to $\frac{\Delta}{2}$ a thin wire since the diameter of the wire equals the cell size.

Interestingly enough $L_{\circledast} = 0$ if the diameter of the wire equals the cell size which in turn means that E_z is zero since the right hand side of (8.9) becomes zero. This is equivalent to the simplest method to include thin wires into the FD-TD scheme, i.e. to put E_z equal to zero for the edges where a wire is located. However, if we stick to the thin-wire model in this case, we get a singular problem, see Section 8.8.

Holland et al. [49] suggests that a correction term is subtracted from L_{\square} to get better agreement to solutions of simple wire problems obtained by the Method of Moments. This correction is obtained by a request to get $L_{\square} = 0$ for $\frac{a}{\Delta} = 0.6$ and

the effect is that of lowering the curve of L_{\square} in Figure 8.6. From all this we can draw the conclusion that it does not matter much whether you use the adjusted L_{\square} or the much simpler L_{\odot} , at least not for quadratic cross-sections.

Without going into any details regarding the thin-wire subcell model of Uma-shankar [103] it is worth noting that putting the wire radius equal to $\frac{\Delta}{e^2}$ ($= 0.135\Delta$) is equivalent to put E_z equal to zero for the edges where a wire is located.

8.3 The wire equations

By making use of the relation between the change in potential and the wire charge density we can write the wire equation in terms of the wire current and the wire voltage. We have $V = \frac{Q}{C}$ where C is the characteristic capacitivity in Farad/meter

$$C = \epsilon\mu \frac{1}{L}. \quad (8.17)$$

We can now reformulate (8.7) and (8.9) and get the system that will be discretized and solved,

$$\begin{aligned} C \frac{\partial V}{\partial t} + \frac{\partial I}{\partial z} &= 0, \\ L \frac{\partial I}{\partial t} + \frac{\partial V}{\partial z} &= E_z - RI. \end{aligned} \quad (8.18)$$

We have here included R which is the resistance per unit length of the wire in Ohm/meter. By changing the boundary condition in (8.1) to $E_z = RI$ and following this through up to (8.9) gives the right hand side $E_z - RI$.

The equation system (8.18) is sometimes called the telegraphers equations and is also called the wave equation formulated as a first-order system. The speed of propagation of the current and voltage is $1/\sqrt{LC} = 1/\sqrt{\epsilon\mu} = c$, i.e. the speed of light.

8.4 Coupling between the wire and the surrounding 3D domain

The electric field surrounding the wire couples to the wire equations as a source term according to (8.18) where the time derivative of the wire current picks up the electric field component aligned with the wire. Coupling back to the 3D domain is accomplished by writing the wire current as a current density. This current density term is treated in the same way as the σE term in Ampère's law, i.e. as a source term for the time derivative of the electric field. This yields a closed form of equations and hence we know how to separate the wire problem from the surrounding field problem, see Figure 8.7 and the system of equations in (8.19).

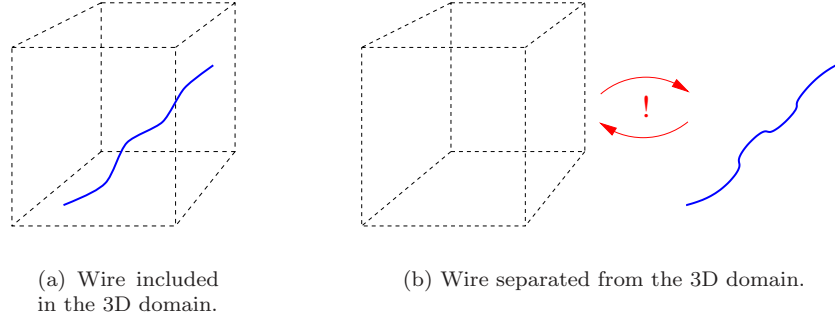


Figure 8.7. Separation of the two problems.

$$\begin{aligned}
 \mu \frac{\partial \mathbf{H}}{\partial t} + \nabla \times \mathbf{E} &= 0, & C \frac{\partial V}{\partial t} + \frac{\partial I}{\partial \xi} &= 0, \\
 \epsilon \frac{\partial \mathbf{E}}{\partial t} - \nabla \times \mathbf{H} &= -\sigma \mathbf{E} - \frac{\mathbf{I}}{A}, & L \frac{\partial I}{\partial t} + \frac{\partial V}{\partial \xi} &= E_\xi - RI,
 \end{aligned} \tag{8.19}$$

where $\frac{\mathbf{I}}{A}$ is the current density associated with the wire current \mathbf{I} , and the cross-section area A is discussed later. (The current density is also denoted \mathbf{J}_w .) The ξ -coordinate is the (local) wire coordinate. If the wire is aligned with the z -axis the ξ -coordinate coincides with the z -coordinate. When we discretize our problem the underlying understanding of these source terms gets more accentuated. In the basic case, where the wire is oriented in a Cartesian direction and the discretized current nodes coincide with the electric components, the coupling between the two equation systems are straightforward. Holland et al. [49] suggests that the wire current is transformed to a current density simply by dividing the current with the cross-section of the Yee cells where the wire is passing. The electric component needed in (8.19) can be picked directly from the 3D grid. Holland et al. [49] also treats the situation where the wire runs parallel to the edges of the Yee cells. In this case a simple bi-linear interpolation is used in order to obtain the electric field at the wire current node. These interpolation weights are re-used when the current density is distributed among the surrounding electric components.

For wires running arbitrarily through the 3D domain we must broaden our view of the governing equations, particularly the mutual source terms.

8.4.1 Distributional contemplation

The wire current is transformed to a current density in (8.19) and if the wire is thin the current density is nonzero only along a one-dimensional curve in \mathbb{R}^3 . This is equivalent to using a one-dimensional Dirac functions in \mathbb{R}^3 for the source term in the Maxwell equations. In the FD-TD method the inclusion of this current density is fairly straight forward if the wire runs along the edges of the FD-TD cell. But

for arbitrarily oriented wires we need to use a careful interpolation technique to couple the FD-TD method and the discretized wire equations.

We start the search for a good interpolation technique by defining a scalar weighting function $\varphi(\mathbf{x}^\xi, \mathbf{x}), \mathbb{R}^{3+3} \rightarrow \mathbb{R}^1$ with the following properties:

1. $\iint_{\Omega} \varphi(\mathbf{x}^\xi, \mathbf{x}) d\mathbf{x} = 1 \quad \forall \mathbf{x}^\xi$ (consistency condition),
2. $\iint_{\Omega} \mathbf{x} \varphi(\mathbf{x}^\xi, \mathbf{x}) d\mathbf{x} = \bar{\mathbf{0}} \quad \forall \mathbf{x}^\xi$ (consistency condition),
3. $\varphi(\mathbf{x}^\xi, \mathbf{x}) \equiv 0 \quad \forall \mathbf{x}^\xi, \mathbf{x} : |\mathbf{x}^\xi - \mathbf{x}| > \rho$ (compact support),
4. $\varphi(\mathbf{x}^\xi, \mathbf{x}) = \varphi(\mathbf{x}, \mathbf{x}^\xi) \quad \forall \mathbf{x}^\xi, \mathbf{x}$ (symmetry),

where \mathbf{x}^ξ is the position of the wire coordinate ξ in 3D space.

The consistency conditions are absolutely fundamental for accuracy reasons. As will be seen later, the weighting function is used for interpolation. The two consistency conditions guarantees that there will be no interpolation error for first-order polynomials. Compact support is necessary in order to avoid large interpolation stencils in the discretization of the equations. Larger stencils are more difficult to implement because surrounding geometries might affect the stencil. Also, the distributed current density should, as much as possible, be concentrated around the corresponding physical current density. A wider stencil distributes the current density over larger domains which degrades the near field accuracy of the model. Finally, symmetry is the key to the stability of the thin-wire model, see Section 8.8.

8.4.2 Distributed coupling between the wire and the surrounding 3D domain

We can now use the weighting function from the previous section to interpolate the surrounding electric field to the wire as well as to distribute the wire currents to the surrounding electric field. Due to the compact support of the weighting function it is enough to distribute and interpolate the quantities within a closed, finite domain. In Figure 8.8 the distribution and interpolation procedure is shown for a compact domain of spherical shape.

We can now define an interpolated value of the ξ component of the electric field. At the wire coordinate ξ we have

$$E_{\xi}^{\text{interp}}(\xi) = \iiint_{\Omega} \varphi(\mathbf{x}^\xi, \mathbf{x}) \mathbf{E}(\mathbf{x}) \cdot \hat{\mathbf{x}}^\xi d\mathbf{x}, \quad (8.20)$$

where $\hat{\mathbf{x}}^\xi$ is the directional vector of the wire. Similarly we define a distributed current density by

$$\mathbf{J}^{\text{distrib}}(\mathbf{x}) = \iiint_{\Omega} \varphi(\mathbf{x}^\xi, \mathbf{x}) I(\mathbf{x}^\xi) \frac{1}{A} \hat{\mathbf{x}}^\xi d\mathbf{x}^\xi \quad (8.21)$$




Figure 8.8. The electric field, at a position on the wire, is obtained by integrating the surrounding electric field multiplied by the weighting function φ (left). The current density at a point in \mathbb{R}^3 is obtained by integrating the current on the wire multiplied by φ and divide by an area (right).

where A is an appropriate area. The volume integral in (8.21) is replaced by a line integral because the wire current only exists along the wire. This yields:

$$\mathbf{J}^{\text{distrib}}(\mathbf{x}) = \int_{\Gamma} \varphi(\mathbf{x}^{\xi}, \mathbf{x}) I(\xi) \frac{1}{A} \hat{\mathbf{x}}^{\xi} d\xi \quad (8.22)$$

The first choice for the area A , is to use the cross-section of an FD-TD cell. For Cartesian oriented wires this is natural, but for arbitrarily oriented wires we have to define the area more carefully, see Section 8.6.2.

8.4.3 Effects of distribution

The consistency in using a distributed current density instead of using the physical current density is illuminated in the following observation.

Assume $\epsilon = 1$ and $\sigma = 0$ for simplicity and study Ampère's law for both the physical case and the distributed case:

$$\begin{aligned} \frac{\partial \mathbf{E}}{\partial t} - \nabla \times \mathbf{H} &= -\mathbf{J}_w, \\ \frac{\partial \tilde{\mathbf{E}}}{\partial t} - \nabla \times \tilde{\mathbf{H}} &= -\tilde{\mathbf{J}}_w, \end{aligned} \quad (8.23)$$

where $\tilde{\mathbf{J}}_w$ is the distributed current density defined via φ in (8.22). We can now study the effect of the distributed current density on $\tilde{\mathbf{E}}$, in a weak sense. We have

$$\frac{\partial}{\partial t} \iiint_{\Omega} (\mathbf{E} - \tilde{\mathbf{E}}) d\mathbf{x} - \iiint_{\Omega} (\nabla \times \mathbf{H} - \nabla \times \tilde{\mathbf{H}}) d\mathbf{x} = - \iiint_{\Omega} (\mathbf{J}_w - \tilde{\mathbf{J}}_w) d\mathbf{x}. \quad (8.24)$$

If we assume the field to be periodic over Ω , the curl of the magnetic fields does not contribute since the curl operator gives derivatives and the integral of a derivative is

zero for periodic functions. That is, $\int_a^b f'(x)dx = f(b) - f(a) = 0$ if $f(x)$ is limited and periodic over the interval. The physical magnetic field is limited because the wire is not infinitely thin and the magnetic field for the distributed case is smoother than in the physical case. Furthermore, if we assume that the domain of support for φ is small enough, such that the current density is constant within this domain, we have

$$\begin{aligned} \iiint_{\Omega} (\mathbf{J}_w - \tilde{\mathbf{J}}_w) d\mathbf{x} &= \iiint_{\Omega} (\mathbf{J}_w - \iiint_{\Omega} \varphi \mathbf{J}_w d\mathbf{x}) d\mathbf{x} \\ &= \iiint_{\Omega} \mathbf{J}_w (1 - \iiint_{\Omega} \varphi d\mathbf{x}) d\mathbf{x} = 0, \end{aligned}$$

where the first consistency condition is utilized. The assumption of a local non-variable current is already introduced in Section 8.2.1 and hence we have, from (8.24),

$$\frac{\partial}{\partial t} \iiint_{\Omega} (\mathbf{E} - \tilde{\mathbf{E}}) d\mathbf{x} = 0.$$

This analysis shows that the difference between the electric fields, using either the distributed current density or the physical current density, is zero in a weak sense.

8.5 Discretization of the wire equations

The equations describing the current and voltage along the wire are discretized by a leap-frog scheme on a staggered grid. In fact the algorithm is more or less identical to the FD-TD algorithm in 1D. This means that the wire model is very well suited to fit into the FD-TD algorithm.

The influence of the surrounding electromagnetic field to the wire is described in Section 8.6 as well as the feedback from the wire to the surrounding electromagnetic field.

The boundary conditions are applied to voltage nodes and hence all current nodes are surrounded by two voltage nodes.

An example of a wire is shown in Figure 8.9 and to simplify the presentation we will assume that the nodes are given in increasing order and that the wire has a start and stop node. However, the implementation can handle wires without start and stop nodes, i.e. circularly connected wires, as well as nodes given in any order.




Figure 8.9. An example of nodes and segments in the xy -plane (left). Corresponding position of voltage and current nodes (right).

8.5.1 Inner nodes

We start by looking at the inner nodes. The voltage- and current-nodes are staggered in both space and time (8.5.1) and the current nodes are placed halfway between the voltage nodes. This gives a centered approximation of $\frac{\partial V}{\partial \xi}$ if approximated at the current nodes. But $\frac{\partial I}{\partial \xi}$ calculated at the voltage nodes will in general not be centered and hence formally only first-order accurate.

The variables are discretized at

$$V|_i^n, \quad i = 1, \dots, N + 1, \quad \text{and} \quad I|_{i+\frac{1}{2}}^{n-\frac{1}{2}}, \quad i = 1, \dots, N, \quad (8.25)$$

where $N + 1$ nodes and N segments are used.

The standard leap-frog scheme is used not only for the Maxwell part of (8.19) but also for the time integration of the wire equations. For wires without resistance we get

$$V_i^{n+1} = V_i^n - \frac{\Delta t}{C} \frac{I_i^{n+\frac{1}{2}} - I_{i-1}^{n+\frac{1}{2}}}{\Delta \xi_I}, \quad (8.26)$$

$$I_{i+\frac{1}{2}}^{n+\frac{1}{2}} = I_{i+\frac{1}{2}}^{n-\frac{1}{2}} - \frac{\Delta t}{L} \left[\frac{V_{i+1}^n - V_i^n}{\Delta \xi_V} - E_\xi \right], \quad (8.27)$$

where $\Delta \xi_I$ is the distance between the current nodes and $\Delta \xi_V$ is the distance between the voltage nodes. For resistive wires (8.27) is replaced by

$$I_{i+\frac{1}{2}}^{n+\frac{1}{2}} = \frac{1 - \frac{R\Delta t}{2L}}{1 + \frac{R\Delta t}{2L}} I_{i+\frac{1}{2}}^{n-\frac{1}{2}} - \frac{1}{1 + \frac{R\Delta t}{2L}} \frac{\Delta t}{L} \left[\frac{V_{i+1}^n - V_i^n}{\Delta \xi_V} - E_\xi \right]. \quad (8.28)$$

The inner points are for the current nodes $1 \leq i \leq N$, and for the voltage nodes $2 \leq i \leq N$, where N is the number of wire segments. Hence, we need boundary

conditions for V_1 and V_{N+1} which is no surprise since these two nodes are the end points of the wire, see Figure 8.9.

8.5.2 Boundary conditions

There are mainly two physical and one numerical boundary condition conceivable. The two physical conditions are either $I = 0$ or $V = 0$ and the numerical boundary condition is the nonreflecting boundary condition.

If a wire-end is free, we assume that there is no current flow through the end point and therefore we enforce the current to be zero at the wire-end. On the other hand, if a wire ends on a large perfectly conducting surface the charge per unit length is zero and therefore we enforce the voltage to be zero at that wire-end. This enforcement is valid only for wires ending on large PECs.

A nonreflecting boundary condition simulates an half-infinite wire and is applied at a wire end point. This is achieved by using a nonreflecting boundary condition, for example the first-order Mur boundary condition.

Boundary conditions for a wire ending on a perfect electric conductor

If $V = 0$ is the desired boundary condition we simply set

$$V_1 = 0, \quad (8.29)$$

$$V_{N+1} = 0. \quad (8.30)$$

Free end boundary condition

If $I = 0$ is the desired boundary condition we plug this implicitly into the time integration of the voltage at the end node by using a mirror value of the current. From (8.26) we have

$$V_1^{n+1} = V_1^n - \frac{\Delta t}{C} \frac{I_1^{n+\frac{1}{2}} - I_0^{n+\frac{1}{2}}}{\Delta \xi_I}, \quad (8.31)$$

$$V_N^{n+1} = V_N^n - \frac{\Delta t}{C} \frac{I_{N+1}^{n+\frac{1}{2}} - I_N^{n+\frac{1}{2}}}{\Delta \xi_I}. \quad (8.32)$$

By putting $I_0 = -I_1$ and $I_{N+1} = -I_N$ we finally get

$$V_1^{n+1} = V_1^n - 2 \frac{\Delta t}{C} \frac{I_1^{n+\frac{1}{2}}}{\Delta \xi_I}, \quad (8.33)$$

$$V_N^{n+1} = V_N^n + 2 \frac{\Delta t}{C} \frac{I_N^{n+\frac{1}{2}}}{\Delta \xi_I}. \quad (8.34)$$

In (8.33) and (8.34) we have used a centered linear extrapolation to achieve $I = 0$ at the wire end. If we instead put $I_0 = -\alpha I_1$, where $\alpha = \frac{1-2\delta}{1+2\delta}$, the effective

length of the wire is increased by $\delta\Delta\xi_I$, $-\frac{1}{2} < \delta < \frac{1}{2}$. This procedure is outlined and verified in [15] for a second-order extrapolation and will not be further addressed in this thesis.

Nonreflecting boundary condition

The nonreflecting boundary condition is useful for long wires. This situation can occur when lightning is simulated and the wire represents the ionized channel where the current is propagating. Another example is for wires running through a PEC surface without Galvanic contact and where the other side of the PEC surface is excluded from the simulation. Since the wire is discretized by a 1D leap-frog scheme the one-way wave equation [31] is perfectly suitable. We get

$$V_1^{n+1} = V_2^n - \frac{c\Delta t - \Delta\xi_V}{c\Delta t + \Delta\xi_V} (V_1^n - V_2^{n+1}), \quad (8.35)$$

$$V_N^{n+1} = V_{N-1}^n - \frac{c\Delta t - \Delta\xi_V}{c\Delta t + \Delta\xi_V} (V_N^n - V_{N-1}^{n+1}). \quad (8.36)$$

8.6 Discrete distribution

The wire equations and the Maxwell equations contains two interaction terms, \mathbf{J}_w and E_ξ , coupling the two system of equations in (8.19) together. For a wire that is discretized along the Cartesian edges of an FD-TD grid there is no need for interpolation in order to communicate electric fields and wire currents, at least not if the current nodes are colocated with the electric components (this case is briefly explained in Section 8.4). But for arbitrarily oriented wires we must use some interpolation technique. In Section 8.4.1 we defined distributional formulations for E_ξ^{interp} and $\mathbf{J}^{\text{distrib}}$, (8.20) and (8.22), that can be used for the interaction between the wire and the surrounding electric field in (8.19).

8.6.1 Introduction to discrete interpolation

We start the discussion in 1D by asking the following question: What is the E_z field in an arbitrary point x ?

The most natural way to determine $E_z(x)$ if x is on a line between two discrete E_z components, is to use the linear interpolation

$$E_z(x^\xi) = \alpha E_z(x_i) + (1 - \alpha) E_z(x_{i+1}) \text{ where } x^\xi = \alpha x_i + (1 - \alpha) x_{i+1}, 0 \leq \alpha \leq 1.$$

This can also be formulated using a discrete weighting function $\varphi(x^\xi, x_i)$, see Figure 8.10, and we have

$$E_z(x^\xi) = \sum_i \varphi(x^\xi, x_i) E_z(x_i).$$




Figure 8.10. A linear discrete weighting function.

It is important that the width of the interpolation function φ is $r = n \Delta x$, where n is an integer, and the amplitude of φ is $1/n$. If this is the case, we inherit the properties of the 3D weighting function discussed in Section 8.4.1 and we have

1. $\sum_i \varphi(x^\xi, x_i) = 1 \quad \forall x^\xi$ (consistency condition),
2. $\sum_i x_i \varphi(x^\xi, x_i) dx = 0 \quad \forall x^\xi$ (consistency condition),
3. $\varphi(x^\xi, x_i) \equiv 0 \quad \forall x^\xi, x_i : |x^\xi - x_i| > \rho$ (compact support),
4. $\varphi(x^\xi, x_i) = \varphi(x_i, x^\xi) \quad \forall x^\xi, x_i$ (symmetry),

In our case we will stick to an interpolation distance equal to $1 \Delta x$ which is the narrowest possible radius. This choice is important in FD-TD simulations since there are so many geometric objects that can obscure the interpolation if the wire runs too close to them. Other commonly used weighting functions include polynomial shaped and trigonometric shaped functions, but for our purpose the linear shaped function works fine. The linear weighting function defined in Figure 8.10 is translation invariant, that is

$$\varphi(x^\xi, x_i) = \varphi(x^\xi - h, x_i - h) \quad \forall h. \quad (8.37)$$

In particular, for $h = x_i$ we have $\varphi(x^\xi, x_i) = \varphi(x^\xi - x_i, 0)$ and from now on we will drop the second argument. Also, due to the symmetry property we have $\varphi(x^\xi - x_i) = \varphi(x_i - x^\xi) = \varphi(|x_i - x^\xi|)$.

Now let us study the 2D case where E_z is sought between four discrete E_z components in the xy-plane. That is, we are looking for the coefficients in the following sum:

$$\begin{aligned} E_z(\mathbf{x}^\xi) = & \alpha_{i,j} E_z|_{i,j} \\ & + \alpha_{i+1,j} E_z|_{i+1,j} \\ & + \alpha_{i,j+1} E_z|_{i,j+1} \\ & + \alpha_{i+1,j+1} E_z|_{i+1,j+1}. \end{aligned}$$

The one-dimensional interpolation technique described above can easily be extended to this situation by simply using the one-dimensional weighting function twice. This gives

$$\begin{aligned}\alpha_{i,j} &= \varphi(x^\xi - x_i) \varphi(y^\xi - y_j) \\ \alpha_{i+1,j} &= \varphi(x^\xi - x_{i+1}) \varphi(y^\xi - y_j) \\ \alpha_{i,j+1} &= \varphi(x^\xi - x_i) \varphi(y^\xi - y_{j+1}) \\ \alpha_{i+1,j+1} &= \varphi(x^\xi - x_{i+1}) \varphi(y^\xi - y_{j+1}),\end{aligned}$$

and is called a bi-linear interpolation. This can also be seen as a tensor product of the two one-dimensional weighting functions and we can write

$$E_z(\mathbf{x}^\xi) = \sum_i \sum_j \varphi(x^\xi - x_i) \varphi(y^\xi - y_j) E_z(x_i, y_j).$$

It is now time to study the general case where E_z is sought in an arbitrary point in \mathbb{R}^3 . In the most general case eight discrete E_z components surrounds the sought E_z and in that case a tri-linear interpolation gives the weights. This is achieved by applying three one-dimensional weighting functions subsequently. This gives

$$E_z(\mathbf{x}^\xi) = \sum_i \sum_j \sum_k \varphi(x^\xi - x_i) \varphi(y^\xi - y_j) \varphi(z^\xi - z_k) E_z(x_i, y_j, z_k), \quad (8.38)$$

where the tensor product of $\varphi(x^\xi - x_i) \varphi(y^\xi - y_j) \varphi(z^\xi - z_k)$ corresponds to the three-dimensional weighting function $\varphi(\mathbf{x}^\xi - \mathbf{x}_i)$ discussed in Section 8.4.

If the desired E_z is located as in the “2D” or “1D” cases above, or even colocated with a discrete E_z point, the tri-linear interpolation is still valid since $\varphi(x_m - x_i) = 1$ if $m=i$ and $\varphi(x_m - x_i) = 0$ if $m \neq i$. Furthermore, it is worth noting that the tensor product of the three one-dimensional linear weighting functions does not destroy the desired properties of consistency, because

$$\iiint_{\Omega} \varphi_x \varphi_y \varphi_z d\mathbf{x} = \int_x \varphi_x dx \int_y \varphi_y dy \int_z \varphi_z dz = 1 \cdot 1 \cdot 1 = 1,$$

where $\varphi_x = \varphi(x^\xi - x_i)$ etc. and

$$\begin{aligned}\iiint_{\Omega} \mathbf{x} \varphi_x \varphi_y \varphi_z d\mathbf{x} &= \iiint_{yz} \varphi_y \varphi_z \int_x x \varphi_x \hat{x} dx dy dz + \iiint_{xz} \varphi_x \varphi_z \int_y y \varphi_y \hat{y} dy dx dz \\ &+ \iiint_{xy} \varphi_x \varphi_y \int_z z \varphi_z \hat{z} dz dx dy = \bar{\mathbf{0}}.\end{aligned}$$

The separation of the integrals are possible because each one-dimensional weighting function is independent of the other two coordinate directions. Replacing the integrals with sums gives the same properties for the discrete case.

8.6.2 Discrete interpolation for the wire problem

The application of the weighting functions can now be used to obtain usable values of tangential electric fields and surface currents in discrete space. For the tangential electric field at the m :th current node we have

$$\begin{aligned}
E_{\xi_m} &= \mathbf{E}^{\text{interp}}(\mathbf{x}_m^\xi) \cdot \hat{\mathbf{x}}_m^\xi = \\
&= \sum_{i,j,k} \varphi(x_m^\xi - i\Delta x - \Delta x/2) \varphi(y_m^\xi - j\Delta y) \varphi(z_m^\xi - k\Delta z) E_x|_{ijk} \hat{\mathbf{x}}_x^{\xi_m} \\
&+ \sum_{i,j,k} \varphi(x_m^\xi - i\Delta x) \varphi(y_m^\xi - j\Delta y - \Delta y/2) \varphi(z_m^\xi - k\Delta z) E_y|_{ijk} \hat{\mathbf{x}}_y^{\xi_m} \quad (8.39) \\
&+ \sum_{i,j,k} \varphi(x_m^\xi - i\Delta x) \varphi(y_m^\xi - j\Delta y) \varphi(z_m^\xi - k\Delta z - \Delta z/2) E_z|_{ijk} \hat{\mathbf{x}}_z^{\xi_m}.
\end{aligned}$$

The summation is broken up into three parts to indicate that the electric components are staggered. Hence the sums gives nonzero contributions for different sets of $\{i, j, k\}$ for different electric components. Equation (8.39) is a direct discretization of the continuously formulated interpolation of (8.20).

Each wire segment m affects the same electric components with the same weights that are involved in (8.39). That is, the current $I(\xi_m)$ is distributed to the electric components within the interpolation radius resulting in a total current density at \mathbf{x} given by

$$\begin{aligned}
J_{E_x}^{\text{distrib}}|_{ijk} &= \sum_m \varphi(i\Delta x + \Delta x/2 - x_m^\xi) \varphi(j\Delta y - y_m^\xi) \varphi(k\Delta z - z_m^\xi) I_{\xi_m} \frac{\Delta \xi_m}{V} \hat{\mathbf{x}}_x^{\xi_m}, \\
J_{E_y}^{\text{distrib}}|_{ijk} &= \sum_m \varphi(i\Delta x - x_m^\xi) \varphi(j\Delta y + \Delta y/2 - y_m^\xi) \varphi(k\Delta z - z_m^\xi) I_{\xi_m} \frac{\Delta \xi_m}{V} \hat{\mathbf{x}}_y^{\xi_m}, \\
J_{E_z}^{\text{distrib}}|_{ijk} &= \sum_m \varphi(i\Delta x - x_m^\xi) \varphi(j\Delta y - y_m^\xi) \varphi(k\Delta z + \Delta z/2 - z_m^\xi) I_{\xi_m} \frac{\Delta \xi_m}{V} \hat{\mathbf{x}}_z^{\xi_m}, \quad (8.40)
\end{aligned}$$

where the ratio between the length $\Delta \xi_m$ and the volume V replaces the area A in (8.22). It is necessary to include the length of the wire segment m to handle wire segments consistently regardless of their length. A motivation for this is clear if we think of two cases. First, assume that the wire is discretized in much shorter segments than the FD-TD edges. In this case the number of segments is large for a given physical wire. Secondly, assume the opposite, that a given wire is discretized in very long segments. In both these cases the electric field should be the same, at least far away from the wire.

Now, assume that the wire current is unitary in both cases. If the current is not scaled by the segment lengths, a lot of short segments will drive the same electric components and hence give rise to non-physically strong current densities. On the other hand, if the segments are long, only a few electric components will be driven by the wire currents. Therefore, it is clear that a length scaling is needed.

To balance the added length dimension a division by volume, instead of area, is needed in (8.22). This leads us to (8.40) where the volume must be constant for all wire segments. The most obvious choice is to take the volume of an FD-TD cell which reduces the ratio $\frac{\Delta\xi_m}{V}$ to the cross-section area for an FD-TD cell in case of Cartesian oriented segments.

8.7 Improved discrete distribution

The discretization of the distributional function φ makes it possible to calculate the electric fields and the current densities needed for the arbitrarily oriented wire segments. However, if we use the tri-linear interpolation technique described in Section 8.6.2 there will be significant discrepancies in the calculated quantities when the segments are slightly moved or rotated (see the results in Section 8.10). In this case the coupling between the wire and the surrounding 3D space is very sensitive to small changes in the wire location. One way to reduce the sensitivity is of course to use a linear weighting function with a radius of $2\Delta x$ instead of Δx , see Figure 8.10. However, numerical experiments show that this does not work very well either.

Inspired by Burke et al. [15] we approach the finding of E_ξ in a two step manner. If we think of the thin wire equations as describing a circular TEM wave guide where we first find the electric field on the shell and then use the averaged shell field for driving the wire we obtain a very favorable interpolation technique. The justification of this approach is primarily based on experimental results.




Figure 8.11. Instead of interpolating the surrounding field directly to the midpoint of the segment, the fields are interpolated to a number of point on a circular shell around the midpoint. These shell point values are then averaged and used at the midpoint.

The shell averaged electric field at the current node is given in vector form by

$$\begin{aligned}\tilde{\mathbf{E}}(\mathbf{x}^\xi) &= \frac{1}{2\pi} \int_0^{2\pi} \sum_{ijk} \varphi(\mathbf{x}^\xi + r\mathcal{A}(\theta)\hat{\mathbf{e}} - \mathbf{x}_{ijk}) \mathbf{E}_{ijk} d\theta \\ &\approx \frac{1}{N} \sum_{n=1}^N \sum_{ijk} \varphi(\mathbf{x}^\xi + r\mathcal{A}(\theta_n)\hat{\mathbf{e}} - \mathbf{x}_{ijk}) \mathbf{E}_{ijk}\end{aligned}\quad (8.41)$$

where r is the shell radius, \mathcal{A} is a rotation matrix describing rotation around the wire-axis and $\hat{\mathbf{e}}$ is an arbitrary unit vector perpendicular to the wire. For reasons of compactness (8.41) is written in vector form. Writing out (8.41) for the three individual electric components would resemble (8.39) but with the averaging sum over the shell points.

This extension of a standard tri-linear interpolation is simple to program once you have a general tri-linear interpolation subroutine. The distribution of the wire current back to the surrounding electric field is done using the same weights that each individual FD-TD component contributes to the respective current nodes, see (8.20) and (8.22). In this way the anti-symmetric structure of the system matrix in Section 8.8 is kept and hence, stability is ensured. The four properties in Section 8.4.1 are also kept. What is left to find is a suitable shell radius and the number of intermediate interpolation points on the shell.

The shell radius is chosen such that the shape of the combined weighting function has a flat center plateau, see Figure 8.12 where the regular bi-linear weighting function and the shell averaged bi-linear weighting function are plotted in the xy -plane. The interpretation of the plotted function of the regular bi-linear case is




Figure 8.12. The regular bi-linear weighting function (left) and the shell averaged bi-linear weighting function (right).

as follows: Assume that the wire node is located in the origin. If there is also an electric component at the origin we pick that electric component with the unit weight (top of the cone). All other components are located outside the domain of support and are hence not involved in the interpolation or distribution. But if the wire is not collocated with an electric component, the regular bi-linear weighting function will give up to four weights for the electric components inside the domain of support. The interpretation of the plotted shell averaged bi-linear weighting function is similar, but here the domain of support is larger which means that there will always be more than one component involved in the coupling process. Furthermore if there is an electric component close to the origin, the exact position is not important at all due to the flatness of the weighting function. The weight for that component will have the plateau value anyway. However, the surrounding electric components will come into play differently if the wire function is moved slightly.

The flatness of the shell averaged function is obtained by using 20 intermediate interpolation points on the circular shell and choosing the shell radius $r = 1/\sqrt{2}$. If r is decreased, the shape will be more conical, and if the radius is increased the shape will resemble a volcano crater. In the latter case centrally placed electric fields will have a too weak influence on the averaged field.

Another commonly used shape function is the cosine bump. Using this shape function instead of the hat function (Figure 8.10) will give a flat plateau for $r = 0.65$. This is thus a more compact combined interpolation function, but numerical experiments indicates that the shell averaged tri-linear weighting function is superior.

8.8 Stability

It is commonly known that both the Maxwell equations and the wire equations are well-posed. But there is no guarantee that a coupling of these two sets of equations preserves the well-posedness. The equations are given in (8.19) but are repeated here for convenience:

$$\begin{aligned} \mu \frac{\partial \mathbf{H}}{\partial t} + \nabla \times \mathbf{E} &= 0, & C \frac{\partial V}{\partial t} + \frac{\partial I}{\partial \xi} &= 0, \\ \epsilon \frac{\partial \mathbf{E}}{\partial t} - \nabla \times \mathbf{H} &= -\sigma \mathbf{E} - \mathbf{J}_w, & L \frac{\partial I}{\partial t} + \frac{\partial V}{\partial \xi} &= E_\xi - RI, \end{aligned} \quad (8.42)$$

where \mathbf{J}_w is the current density given by the current I divided by an area, see Section 8.6.2. Studying the well-posedness of the coupled equations is very complicated because they have different dimensionality. The Maxwell equations are defined in three space dimensions, and the wire equations are one-dimensional in space. Therefore, let us study the equivalent property of stability for the discretized equations. But also in the discrete case a general stability study is very difficult. Hence, we will have to make a few assumptions: Assume that the wire is straight, oriented in z-direction, and infinitely long. Also assume that there is no variation

of any variables in the z -direction and that the electric field only has an E_z component. Furthermore conductivity and resistance per unit length is omitted. Now we can reduce (8.42) to the 2D TM _{z} mode, see (3.8),

$$\begin{aligned} \mu \frac{\partial H_x}{\partial t} &= -\frac{\partial E_z}{\partial y}, & C \frac{\partial V}{\partial t} + \overbrace{\frac{\partial I}{\partial z}}^0 &= 0, \\ \mu \frac{\partial H_y}{\partial t} &= \frac{\partial E_z}{\partial x}, & L \frac{\partial I}{\partial t} + \underbrace{\frac{\partial V}{\partial z}}_0 &= E_z, \\ \epsilon \frac{\partial E_z}{\partial t} &= \frac{\partial H_y}{\partial x} - \frac{\partial H_x}{\partial y} - J_z, \end{aligned} \quad (8.43)$$

and the wire equations reduces to an ordinary differential equation for I .

A necessary condition for stability of the fully discrete problem is that the equivalent semi-discrete problem, written on matrix form in (8.44), have eigenvalues in the left half plane. For the leap-frog integrator to be stable, the eigenvalues must be purely imaginary.

$$\frac{\partial}{\partial t} \begin{pmatrix} H_x \\ H_y \\ E_z \\ I \end{pmatrix} = \begin{pmatrix} \mathbf{0} & \mathbf{0} & \mathcal{A} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathcal{B} & \mathbf{0} \\ \mathcal{C} & \mathcal{D} & \mathbf{0} & \mathcal{F} \\ \mathbf{0} & \mathbf{0} & \mathcal{G} & \mathbf{0} \end{pmatrix} \begin{pmatrix} H_x \\ H_y \\ E_z \\ I \end{pmatrix}, \quad (8.44)$$

where the operators \mathcal{A} and \mathcal{B} take care of the spatial discretization of E_z , and \mathcal{C} and \mathcal{D} take care of the spatial discretization of H_x and H_y . The operator \mathcal{F} describes how the wire current affects the electric field and the operator \mathcal{G} describes how the electric field affects the wire current.

If we introduce a scaling of the variables the symmetry of the operators are obvious and the numerical solution of the eigenvalue problem becomes more robust. Using the following scaling:

$$h_x = \sqrt{\mu} H_x, \quad h_y = \sqrt{\mu} H_y, \quad e_z = \sqrt{\epsilon} E_z, \quad i = \sqrt{\mu} \sqrt{g(a)} I, \quad (8.45)$$

and substituting this into (8.44) only changes the magnitude of the operators $\mathcal{A} - \mathcal{G}$. The scaling of the current is $1/\sqrt{L}$ where we split L in two parts, the permeability, and $g(a)$. From (8.16) we have

$$g(a) = \frac{1}{2\pi} \ln\left(\frac{\Delta}{2a}\right) \text{for } L_{\otimes}. \quad (8.46)$$

After the substitution, all the operators in (8.44) have the speed of light as a common factor, which can be extracted. If we rename our new variables and operators, and choose the old names, equation (8.44) is valid again, but scaled.

Let us study an example. Assume that the wire is located in the center of a 6×6 cells FD-TD TM _{z} -grid, see Figure 8.13. Use a Dirichlet boundary condition




Figure 8.13. The position of the discretized variables. PMC boundary condition is applied to the tangential magnetic components along the dashed lines.

for the tangential magnetic components in the outermost layer of cells, i.e. a PMC boundary condition. This gives us 5×4 H_x components, 4×5 H_y components, 5×5 E_z components, and one current component. The total number of unknowns are thus 66. Without loss of generality we assume that $\Delta x = \Delta y = 1$. If this is not the case the scaling of the current must take Δx and Δy into account as well.

If we order the unknown lexicographical, that is, starting from the lower left corner going to the right and then next row, the center E_z component is E_z^{13} . The operators $\mathcal{A} - \mathcal{D}$ in (8.44) have a simple structure and are mutually anti-symmetric. That is,

$$\mathcal{C}^T = -\mathcal{A}, \quad (8.47)$$

$$\mathcal{D}^T = -\mathcal{B}, \quad (8.48)$$

$$\mathcal{G}^T = -\mathcal{F}, \quad (8.49)$$

where \mathcal{F} is a 25×1 matrix, and \mathcal{G} is a 1×25 matrix. The only nonzero components of \mathcal{F} and \mathcal{G} are

$$\mathcal{F}(13, 1) = -\frac{c}{\sqrt{(g)}}, \quad \mathcal{G}(1, 13) = \frac{c}{\sqrt{(g)}}, \quad (8.50)$$

if the wire is placed in the center of the domain.

The structure of the total matrix in (8.44) is plotted in Figure 8.14. The different sections are mutually anti-symmetric and hence the semi-discrete problem is described by an anti-symmetric matrix, which has purely imaginary eigenvalues.




Figure 8.14. The structure of the matrix in (8.44). The different block are the operators \mathcal{A} - \mathcal{G} and the nonzero elements (plotted) are minus one or plus one.

To prove stability for the fully discrete case, using the leap-frog time integrator, we study the magnitude of the imaginary eigenvalues. A sufficient condition for our example is that the largest eigenvalues are contained within the stability domain of leap-frog.

If we calculate the largest eigenvalue to the problem described in Figure 8.13 for different number of cells we will find that the eigenvalues are slightly larger than for a problem without a wire, see Table 8.8. Therefore, the largest time step allowed for the wire problem is slightly smaller than what is allowed for the corresponding stand alone problem. How much the time step must be reduced depends on the wire radius. If the wire radius is increased the maximum time step must be decreased. From Table 8.8 we have that the time step must be reduced by a factor 1.06 for a wire with relative radius of 0.1. For a wire with relative radius of 0.25 the time step must be decreased by a factor 1.30. But as long as we use a sufficiently small time step we have stability.

Problem-size	no wire	Wire radius			
		a=1.7410 ⁻⁶	a=0.00093	a=0.0216	a=0.1
	Maximum eigenvalue				
n-1	λ_{empty}	$\lambda_{g=2}$	$\lambda_{g=1}$	$\lambda_{g=0.5}$	$\lambda_{g=0.256}$
5	0.9511	0.9570	0.9648	0.9872	1.0575
10	0.9877	0.9892	0.9914	1.0013	1.0593
15	0.9945	0.9952	0.9964	1.0031	1.0593
20	0.9969	0.9973	0.9980	1.0035	1.0593
25	0.9980	0.9983	0.9988	1.0036	1.0593
30	0.9986	0.9988	0.9992	1.0036	1.0593
35	0.9990	0.9991	0.9994	1.0036	1.0593
40	0.9992	0.9994	0.9996	1.0036	1.0593
45	0.9994	0.9995	0.9997	1.0036	1.0593
50	0.9995	0.9996	0.9997	1.0036	1.0593

Table 8.1. Maximum eigenvalue for the semi-discrete case for different problem sizes, $n_x = n_y = n$, and different wire radius a . g is the scaling factor defined in (8.46).

The numerical analysis performed so far is for a wire collocated with the center electric component. The analysis can also be done for wires that are located arbitrarily in the grid. For a wire located in the center of a cell, the operators \mathcal{F} and \mathcal{G} will have four nonzero entries, at least if a bi-linear interpolation function is used. The values of the matrix entries is one fourth of (8.50). For other wire position the fractions between the four nodes are different. Numerical experiments on the case where the current node is offseted $0.25\Delta x$ and $0.25\Delta y$ resembles the results in Table 8.8 even though the wire radius does not have the same impact on the eigenvalues as in the collocated wire case. For the problem size $n - 1 = 50$ the maximum eigenvalue is only 1.0009 for a wire radius of $0.25\Delta x$.

An important question is how general this stability study is. Does it apply to 3D using a wire with finite length? First of all we realize that the semi-discrete 3D problem without a wire can be described by an anti-symmetric matrix because the leap-frog time integrator is stable as long as the time step is limited by (4.10). Secondly, even though the size of \mathcal{F} and \mathcal{G} becomes larger (one column/row for each current node), we still have $\mathcal{F}^T = -\mathcal{G}$ because the coupling from the field to the wire is identical to the opposite coupling, except for a change in sign. We will have to scale and include the voltage as well but all the operators will be mutually anti-symmetric and hence the stability is ensured as long as the time step is chosen according to the largest eigenvalues of the semi-discrete problem. However, for a realistic 3D case, it is very cumbersome to calculate the largest eigenvalue because the semi-discrete matrix becomes very large, and consequently, a practical limit on the time step is in general very difficult to obtain.

8.9 Comments on the physical assumptions

There are reasons to believe that the thin-wire model is accurate in the case of long straight wires. But bearing in mind the introductory consideration of infinite long cylinders it is dubious if the wire model always mimics the physics. What happens close to wire-ends, where the wire bends abruptly, or if the wire runs close to material interfaces?

At a wire-end the assumption of an $\frac{1}{r}$ dependency of the electric field on wire charge, described by Biot-Savart's law (8.6), is not valid. This effect is studied in [26] for the wire model of Umashankar et al. [103] and a significant effect is shown.

Whether bending wires are poorly modeled is most of all a question of wire radius. The coupling between the surrounding electromagnetic field and the wire is stronger the thicker the wire is. Therefore, for really thin wires one can argue that most energy will propagate in the wire and only a small fraction will be emitted before the bend and collected after the bend, hence the physics is not severely violated by the assumption for thin wires.




Figure 8.15. Conducted and radiated power close to a bending wire.

Close to material interfaces, the field is not decreasing like $\frac{1}{r}$ as assumed in equation (8.5) and equation (8.6). We can perhaps recover the assumption and modify the corresponding equations by considering a mirrored problem in the case of wires running close to perfect electric conductors, see Figure 8.16. For general materials, however, the method of mirroring becomes unpractical, particularly if the wire is not running parallel to the geometry.




Figure 8.16. The field around a cylinder over an infinite ground plane can be found using a mirror-cylinder.

8.10 Results

In this section the thin-wire subcell model is used for modeling two antennas, a linear dipole antenna and a circular loop antenna. The dipole antenna is run in both receiving and transmitting mode.

The coupling from the surrounding 3D domain to the thin wire is most critical for the receiving antenna whereas the the coupling from the thin wire to the 3D domain is most important for the transmitting antenna.

The loop antenna demonstrates the capability of arbitrary orientation as well as curved wires.

Two interpolation techniques are examined, the standard tri-linear interpolation which is a natural extension to the bi-linear interpolation proposed by Holland et al. [49], and the shell averaged tri-linear interpolation from Section 8.6.

The new shell averaged tri-linear interpolation technique gives very good results, which is shown in the following experiments. But the standard tri-linear interpolation is incapable of modeling arbitrarily oriented thin wires correctly.

In preparation for Section 8.10.1 and 8.10.2 some fundamental properties of a dipole antenna are given here (more details are found in Balanis [7]). An antenna is a device that interacts with its surrounding media and for certain frequencies the interaction is particularly strong. These frequencies corresponds to the resonant modes of the wire and for a linear dipole the first resonant modes are plotted in Figure 8.17. If the antenna length in Figure 8.17 is 41 meter the corresponding




Figure 8.17. The first five resonant modes on a dipole antenna. The horizontal line is the wire, the solid curve is the current and the dashed curve is the voltage.

wavelengths for the eigenmodes are given in Table 8.2.

mode	λ [m]	f [MHz]	mode	λ [m]	f [MHz]
1	82.0	3.6585	5	16.4	18.2927
2	41.0	7.3171	6	13.6667	21.9512
3	27.3333	10.9756	7	11.7143	25.6098
4	20.5	14.6341	8	10.25	29.2683

Table 8.2. The first eight resonant modes for a 41 m long wire, see Figure 8.17. The frequency is based on the free space wave propagation speed.

8.10.1 Receiving dipole antenna

A plane wave impinging a dipole antenna will induce currents on the antenna. The currents will propagate along the antenna to the end point and reflect back again. This process will continue until the energy induced by the exciting field has been radiated out to the surrounding medium. For some frequencies this process is particularly strong due to resonances. The resonant frequencies are easier to pick up by the antenna and hence the corresponding currents are stronger.

In the following numerical experiment a thin wire with a length of 41 meter is embedded in an FD-TD grid with $\Delta x = \Delta y = \Delta z = 1$ m. The wire radius is 10 mm and the wire runs in the x -direction discretized in one meter segments. The direction of the incoming pulse is 26.6 degrees ($\text{atan}(\frac{1}{2})$) from broadside and a horizontal polarization is used in order to excite the wire. By studying the radiation

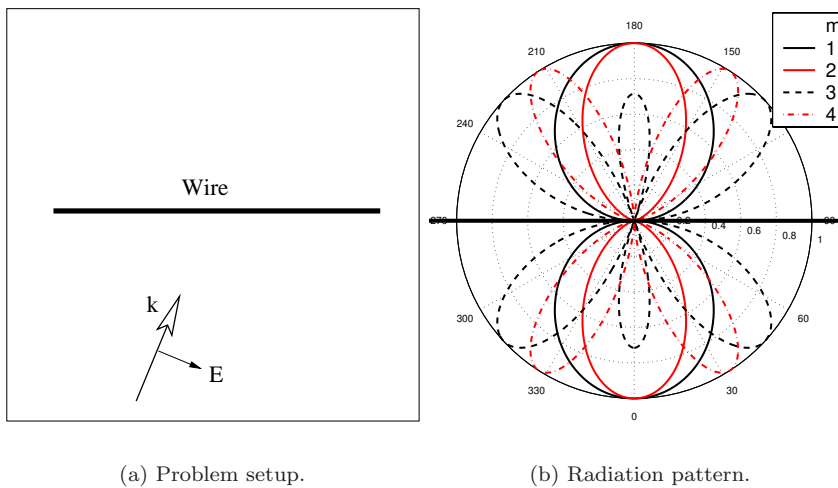


Figure 8.18. A plane wave impinges the wire at an angle of 26.6 degrees with a plane polarization indicated to the left. To the right the radiation pattern (linear) for the first four resonant modes.

pattern for a transmitting dipole antenna we find the angle 26.6 degrees for which the first ten modes does not vanish. Reciprocity gives us that the same angle will excite the modes if the impinging plane wave contains the corresponding frequency components.

The pulse shape is a differentiated Gaussian pulse with $t_w = 2.0 \times 10^{-8}$ i.e. the -3 dB bandwidth is 3.6 – 22 MHz. We use the time step $\Delta t = 0.5\Delta x/c_0$. Current node 21 is the midpoint of the wire.

In the basic setup the current nodes coincide with E_x components and the wire is then moved in four different ways.

In the first case the wire is moved transversally as in Figure 8.19a using $0.25\Delta y$ (Ia_25), $0.5\Delta y$ (Ia_50) and $0.5\Delta y + 0.5\Delta z$ (Ia_55). In the second case the wire is moved longitudinal as in Figure 8.19b using $0.25\Delta x$ (Ib_25) and $0.5\Delta x$ (Ib_50). In the third case the wire is rotated in the xy-plane, Figure 8.19c, by 26.6 degrees (Ic_27) and 45 degrees (Ic_45). In the fourth case the wire is rotated as in Figure 8.19d from the direction $[1\ 0\ 0]$ to $[1\ 0.5\ 0.5]$ (Id_24) and to $[1\ 1\ 1]$ (Id_35).

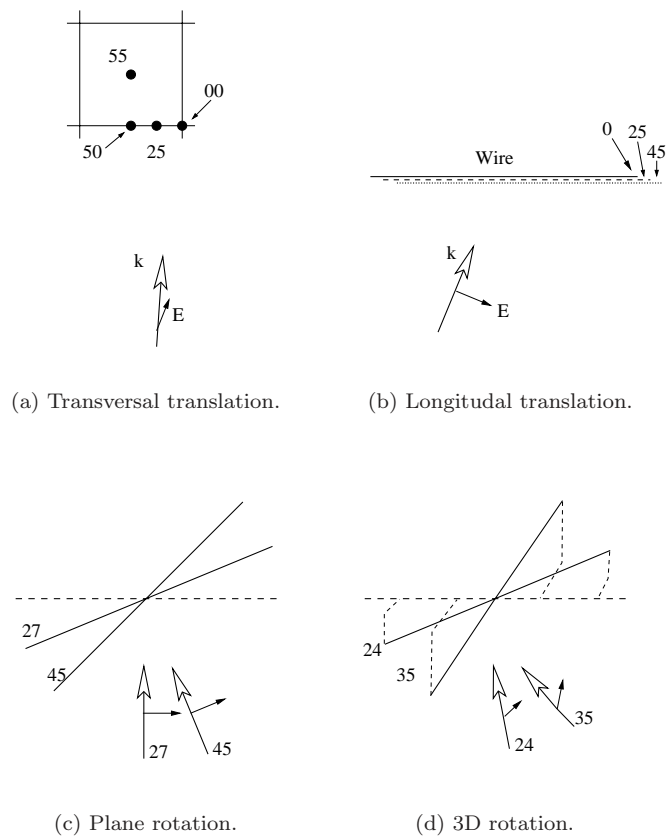


Figure 8.19. The four different test cases for moving the wire. Case a) is seen from the side, and in case b) the wire is moved in the length direction.

In all these cases the incoming wave is designed such that the time when the wave reaches the center point of the wire is the same, and hence a direct comparison of the currents and voltages can be performed. The responses of the different cases are shown in Figure 8.20–8.23.

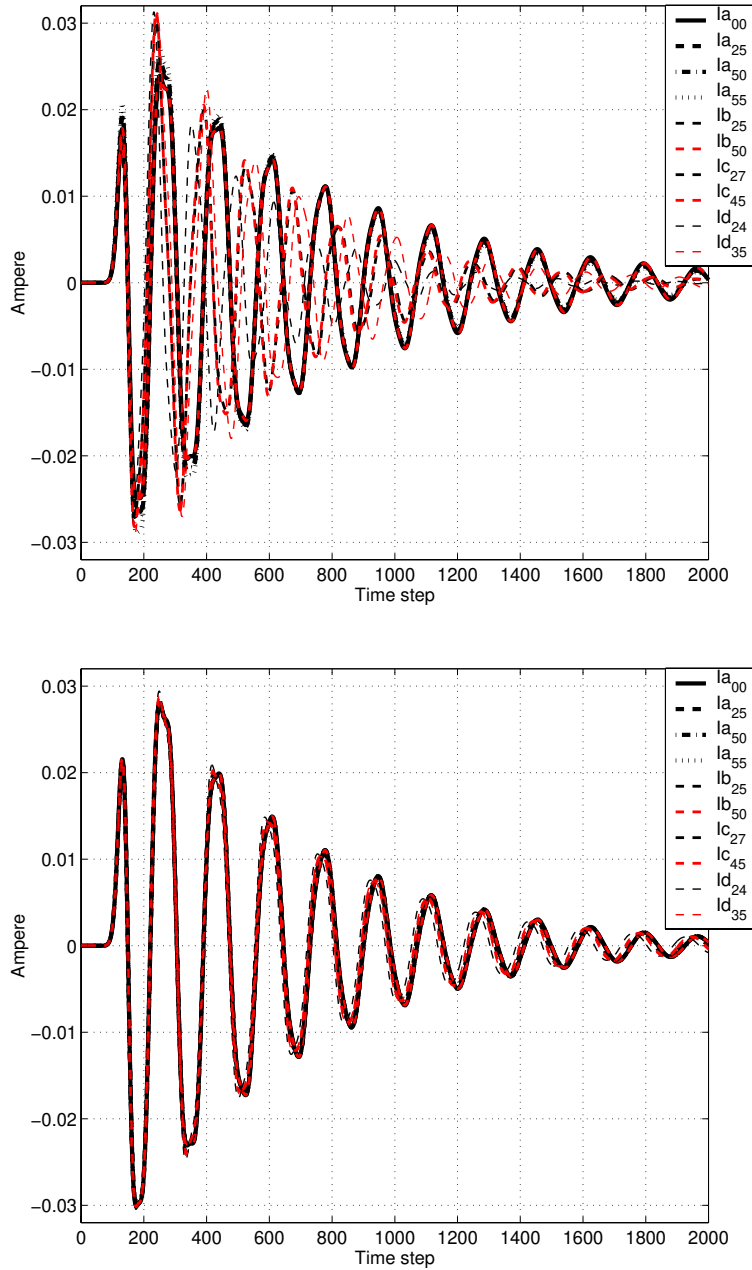


Figure 8.20. The midpoint current for the standard tri-linear interpolation (top) and the shell averaged tri-linear interpolation (bottom). Close ups are given in the following two figures.

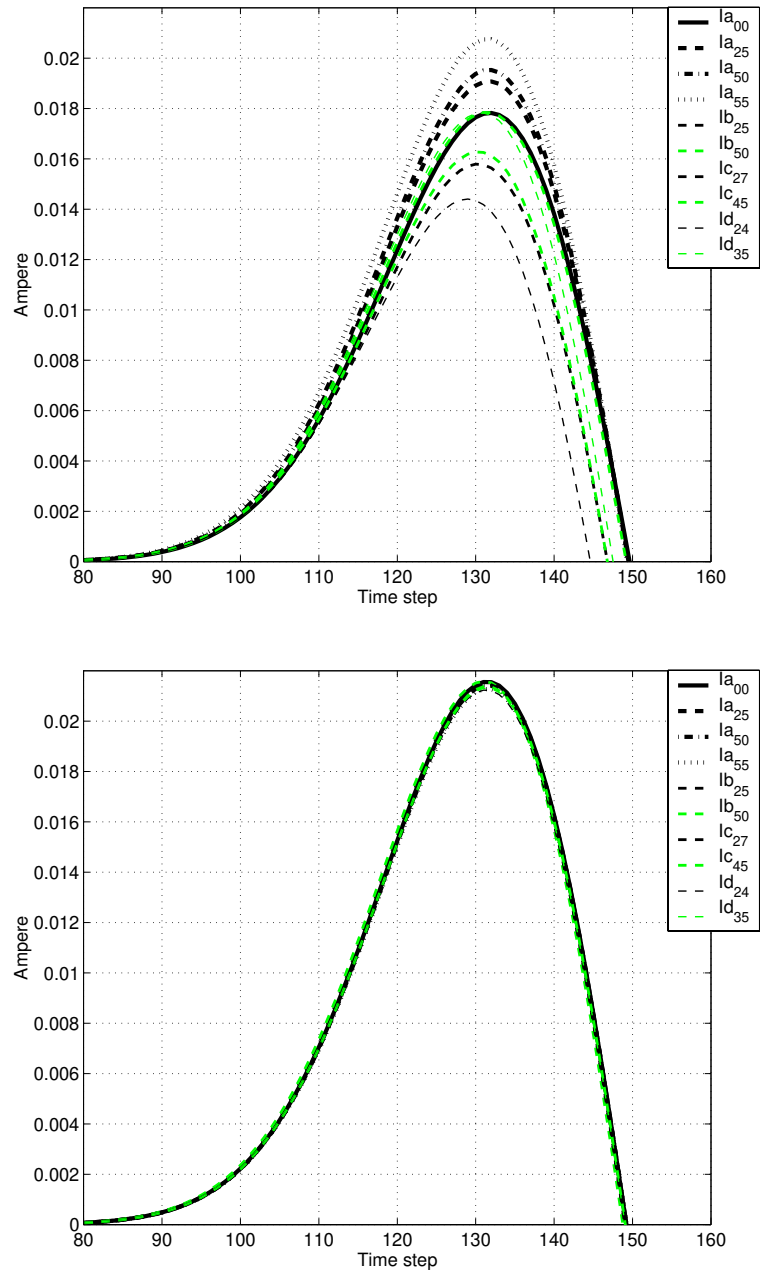


Figure 8.21. Detail of the midpoint current for the standard tri-linear interpolation (top) and the shell averaged tri-linear interpolation (bottom).

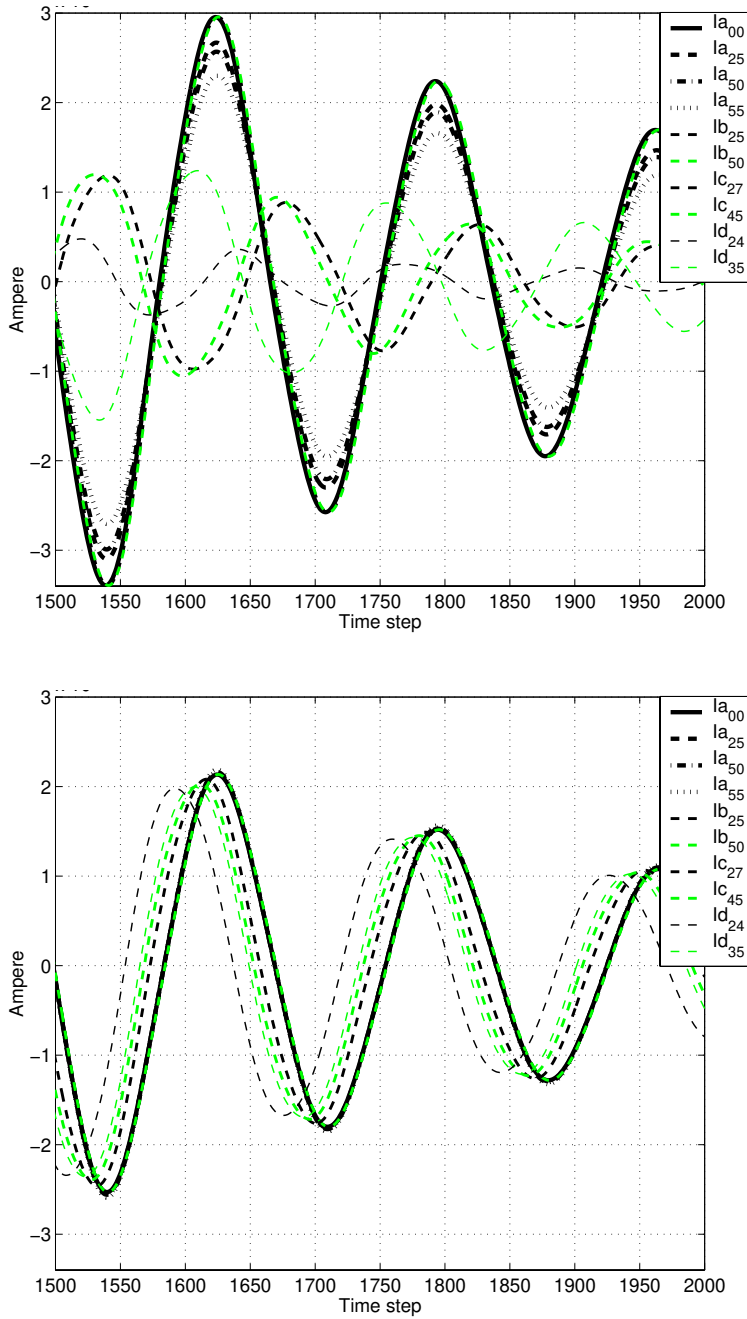


Figure 8.22. Detail of the midpoint current for the standard tri-linear interpolation (top) and the shell averaged tri-linear interpolation (bottom).

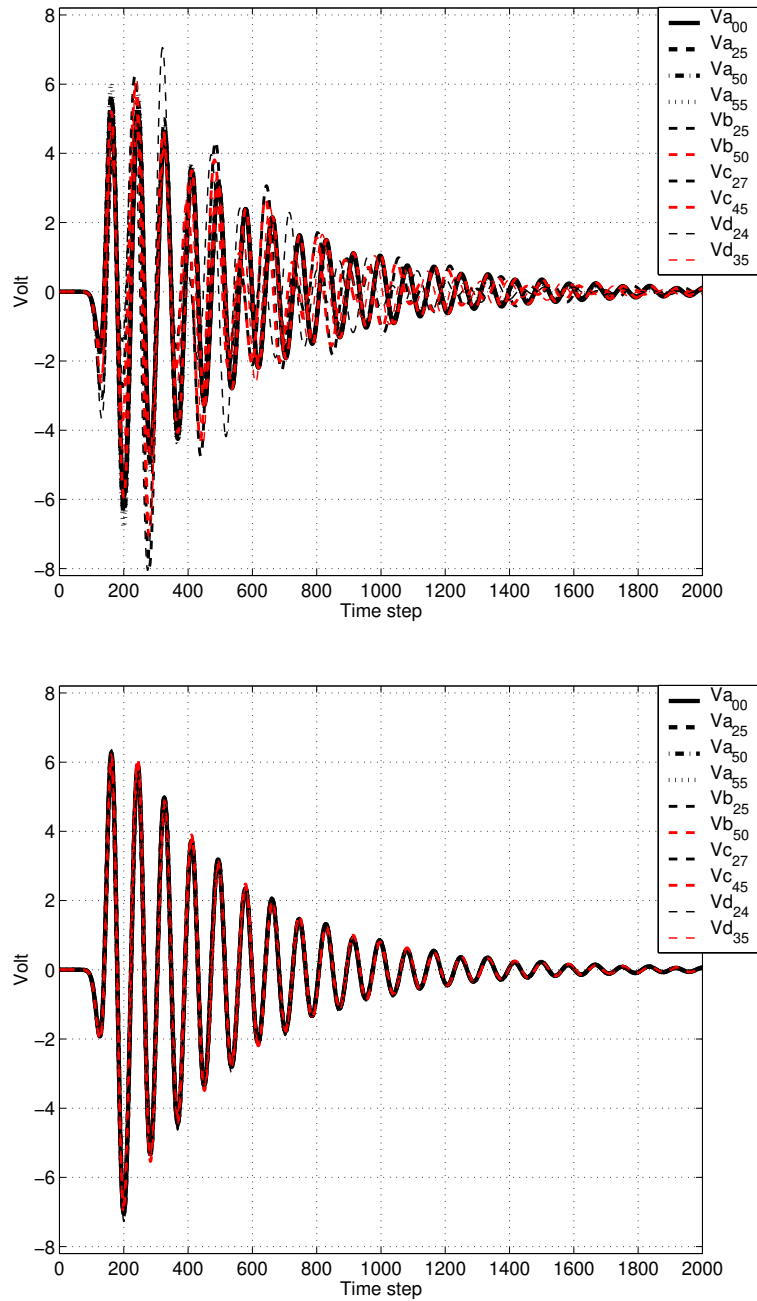


Figure 8.23. The (almost) midpoint voltage V_{22} for the standard tri-linear interpolation (top) and the shell averaged tri-linear interpolation (bottom).

It is tempting to believe that the responses shown in Figure 8.20 and Figure 8.23 is described by damped harmonics, at least after a while when the initial excitation have settled. Assuming that the response is described by a sum of damped oscillations we can find them using Prony's method. This method is described in the literature, see for example [98], but it is necessary to use the method with some care since spurious modes might occult the expansion.

Assume that the current response can be written

$$I_{21}(t) = \sum_{l=1}^p C_l e^{(-\gamma_l t + i2\pi f_l t)}, \quad (8.51)$$

where $Abs(C_l)$ is the initial amplitude, $Arg(C_l)$ is the phase, γ constitutes the damping, and f is the frequency of the oscillation.

If we run Prony's method and extract the most dominant oscillation we can summarize the responses for the two interpolation techniques in a table for direct comparison. We get:

Case	Standard tri-linear interpolation.			Shell averaged interpolation.		
	f [MHz]	γ [$\frac{1}{\mu s}$]	ampl. [mA]	f [MHz]	γ [$\frac{1}{\mu s}$]	ampl. [mA]
Ia_00	3.527	0.973	20.02	3.519	1.179	25.49
Ia_25	3.524	1.041	21.77	3.519	1.175	25.39
Ia_50	3.523	1.066	22.42	3.519	1.171	25.28
Ia_55	3.520	1.134	24.24	3.519	1.163	25.06
Ib_25	3.523	0.972	19.99	3.516	1.178	25.46
Ib_50	3.521	0.971	19.97	3.515	1.178	25.45
Ic_27	4.169	1.437	24.42	3.543	1.193	25.67
Ic_45	4.191	1.444	25.34	3.554	1.201	25.71
Id_24	4.672	1.890	28.52	3.594	1.223	26.03
Id_35	3.972	1.379	25.56	3.564	1.204	25.72

Table 8.3. The dominant current mode for a 41 m long wire using standard tri-linear interpolation and shell averaged tri-linear interpolation.

The relative changes between the different wire cases is shown for both interpolation techniques in Table 8.4. The nominal values are chosen by inspection of the figures in Table 8.3.

Case	Relative difference in percent, standard tri-linear interpolation.			Relative difference in percent, shell averaged interpolation.		
	Nominal values,			Nominal values,		
	3.52 [MHz]	1.18 [$\frac{1}{\mu s}$]	25.4 [mA]	3.52 [MHz]	1.18 [$\frac{1}{\mu s}$]	25.4 [mA]
Ia_00	0.2	-21.2	-26.9	.	-0.1	0.4
Ia_25	0.1	-13.4	-16.7	.	-0.4	.
Ia_50	0.1	-10.7	-13.3	.	-0.7	-0.5
Ia_55	.	-4.1	-4.8	.	-1.4	-1.4
Ib_25	0.1	-21.4	-27.1	-0.1	-0.2	0.3
Ib_50	.	-21.5	-27.2	-0.1	-0.2	0.2
Ic_27	15.6	17.9	-4.0	0.6	1.1	1.1
Ic_45	16.0	18.3	-0.2	1.0	1.8	1.2
Id_24	24.7	37.6	10.9	2.1	3.5	2.4
Id_35	11.4	14.4	0.6	1.2	2.0	1.2

Table 8.4. The relative differences in the dominant current mode for a 41 m long wire using standard tri-linear interpolation and shell averaged tri-linear interpolation.

It is very clear that standard tri-linear interpolation does not give consistent results for different discretizations of the same physical setup. For pure translation (Ia_00 to Ib_50) the frequency is accurate, but the damping and amplitude varies several ten's of percent. For the rotated wires (Ic_27 to Id_35) also the frequency varies very much. The shell averaged interpolation technique gives a variation of only a few percent for the worst cases.

8.10.2 Transmitting dipole antenna

In this section we study the same dipole antenna as in the previous section but in transmitting mode instead of receiving mode. The new shell interpolated wire model showed very consistent results in receiving mode and it is now time to ask the obvious question whether the model gives consistent and accurate results also for transmitting antennas?

A transmitting dipole antenna is characterized by the input impedances and input admittances. Burke et al.[15] present very good results using their method for transversally moved wires. But a more general capability to tilt a wire with respect to the mesh is inquired in the conclusion of[15]. In this section such results are presented using the new shell interpolated model.

The transmitting dipole antenna is excited at the midpoint of the antenna where also the input impedance is registered. We excite the center current segment using a voltage source which is added to the right hand side of (8.27);

$$I_{i+\frac{1}{2}}^{n+\frac{1}{2}} = I_{i+\frac{1}{2}}^{n-\frac{1}{2}} - \frac{\Delta t}{L} \left[\frac{V_{i+1}^n - V_i^n}{\Delta \xi_V} + E_\xi \right] + \frac{\Delta t}{L \Delta \xi_V} V^{\text{exc.}}(t_n).$$

If we register the current at the mid segment we can calculate the input impedance and input admittance for the dipole antenna. The impedance is calculated as the ratio between the input voltage and the current in the frequency domain but we have to compensate for the time-staggering between the excited voltage (t_n) and the registered current ($t_{n+\frac{1}{2}}$). If we use one of the nominal Fourier transformed quantities we can choose to compensate the other by using either $\hat{V}^{\text{exc.}}(f) = \text{fft}(V^{\text{exc.}}(t_{n+\frac{1}{2}}))$, $\hat{I}_{21}(f) = \text{fft}(I_{21}^{n+\frac{1}{2}} e^{-i2\pi f dt/2})$ or $\hat{I}_{21}(f) = \text{fft}(I_{21}^{n-\frac{1}{2}} + I_{21}^{n+\frac{1}{2}}) * 0.5$. However, the differences in the results are very small regardless of the choice.

The impedance is given by

$$\hat{Z}_{21}(f) = \frac{\hat{V}^{\text{exc.}}(f)}{\hat{I}_{21}(f)}, \quad (8.52)$$

and the admittance

$$\hat{Y}_{21}(f) = \frac{1}{\hat{Z}_{21}(f)}. \quad (8.53)$$

We can split the impedance and admittance into the real and imaginary parts. The real and imaginary components are called resistance and reactance for the impedance, and conductance and susceptance for the admittance.

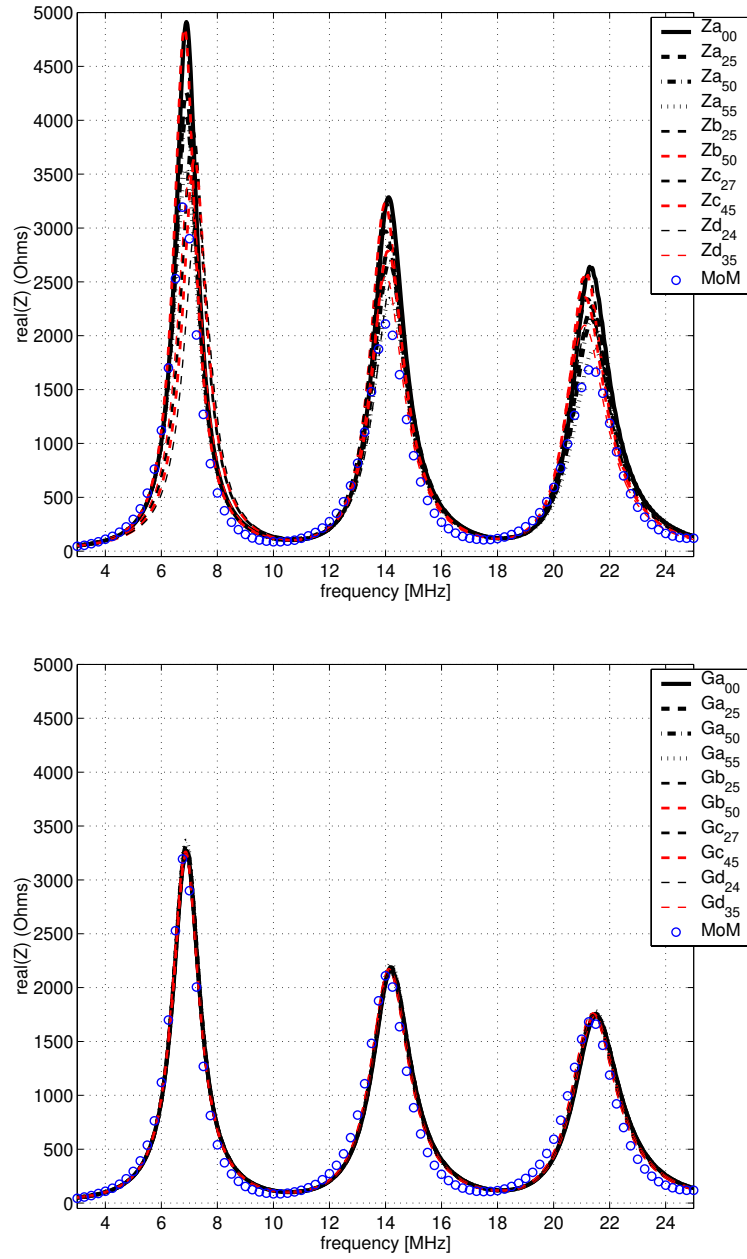


Figure 8.24. The input resistance for the standard tri-linear interpolation (top) and the shell averaged tri-linear interpolation (bottom). A Method of Moments reference solution is also plotted.

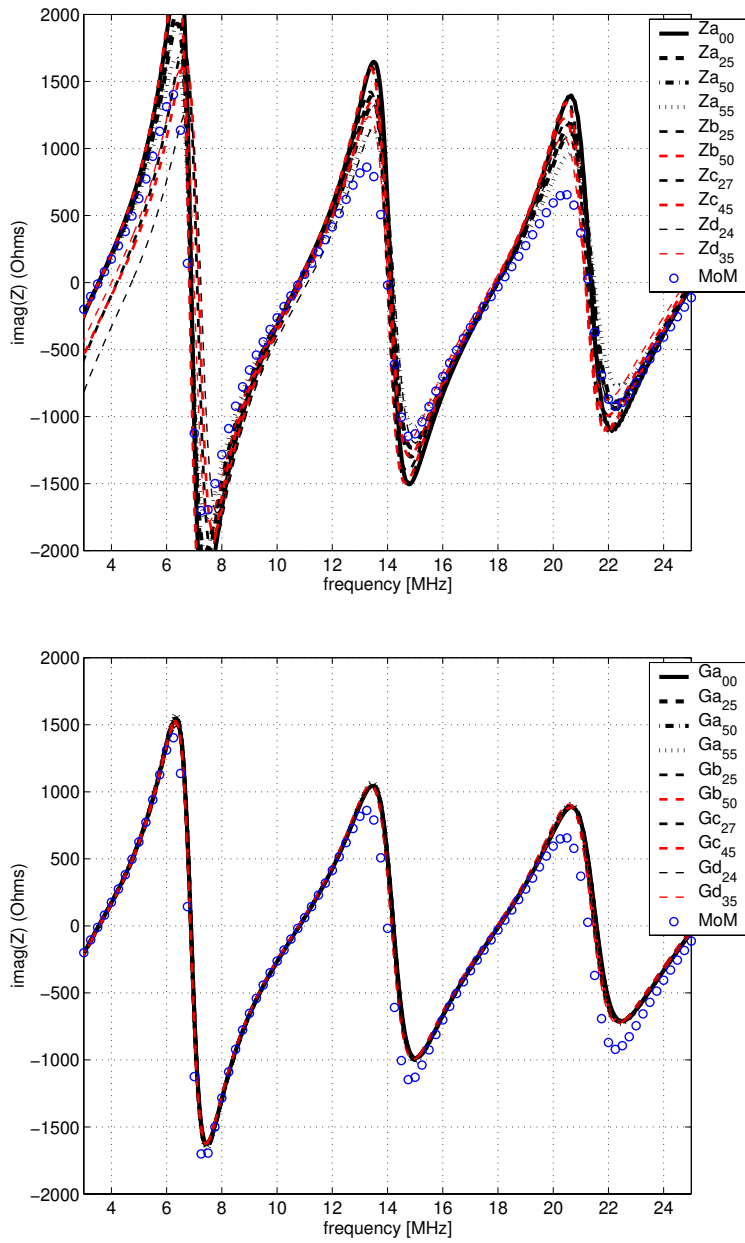


Figure 8.25. The input reactance for the standard tri-linear interpolation (top) and the shell averaged tri-linear interpolation (bottom). A Method of Moments reference solution is also plotted.

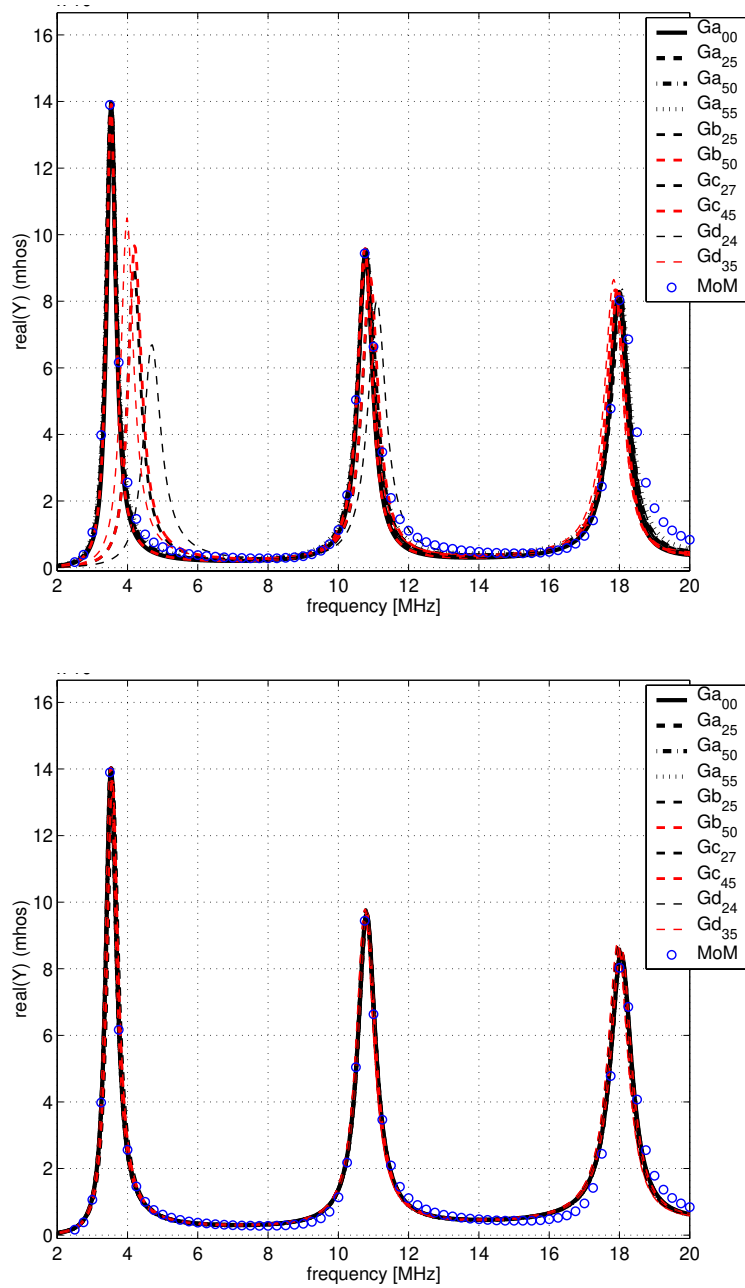


Figure 8.26. The input conductance for the standard tri-linear interpolation (top) and the shell averaged tri-linear interpolation (bottom). A Method of Moments reference solution is also plotted.

The impedance and admittance for the different wire cases (see Figure 8.19) are plotted in Figure 8.24 – 8.26 where we have some very clear resonances. These correspond to the resonance modes shown in Figure 8.17 where the first, third and fifth mode have nonzero currents but zero voltage whereas the second and fourth mode have nonzero voltages and zero currents. Hence, we can expect peaks in the impedances for the modes where the currents are zero, i.e. the even modes. The odd modes will give peaks in the admittances.

The results are compared to a Method of Moments solution obtained by NEC-3 and we see a good agreement for the shell interpolated wire model. But the standard tri-linear interpolation technique gives very poor results. The impedance for the first half wavelength resonance should be $73 + i42.5 \Omega$ [7] and the reactance goes from being capacitive to inductive very rapidly in the vicinity of the resonance frequency (see Figure 8.25). Therefore we expect the impedance to be fairly close to 73Ω at the frequency where the reactance goes to zero around the half wavelength resonance. This also corresponds to where the admittance has its first peak. The input impedance for the different wire cases are given in Table 8.5.

The impedances obtained for the standard tri-linear interpolation technique are good as long as the wire is oriented along a Cartesian direction. But if we rotate the wire we get very bad impedances as well as too high resonance frequencies. The shell averaged interpolation technique gives good results in all cases. Interestingly enough the impedance and resonance frequencies are closer to analytic values (73Ω and $3.65 MHz$) for the rotated wires. The reason for this is not clearly understood, but we note that a reference solution obtained by NEC-3 gives $3.53 MHz$ and 72.2Ω which is also slightly below expected analytic values. A more careful study of the influence of the wire radius might give some insight. The thicker the wire, the tighter the coupling between the 3D field problem and the 1D wire problem becomes, and the excitation and radiation process becomes stronger. This in turn will lead to a higher damping term in the expansion in (8.51) which will lead to a shift in frequency, see Section 8.10.4.

Case	Standard tri-linear interpolation.		Shell averaged interpolation.	
	f [MHz]	real(Z) [Ω]	f [MHz]	real(Z) [Ω]
Ia_00	3.54	71.9	3.54	71.6
Ia_25	3.54	71.8	3.54	71.6
Ia_50	3.54	71.7	3.54	71.6
Ia_55	3.54	71.6	3.54	71.6
Ib_25	3.54	71.8	3.54	71.5
Ib_50	3.54	71.8	3.54	71.5
Ic_27	4.20	111.3	3.57	72.3
Ic_45	4.22	104.0	3.58	73.2
Id_24	4.73	151.2	3.62	74.1
Id_35	4.00	95.8	3.59	73.5

Table 8.5. The frequency and resistance at half wavelength resonance for a 41 m long transmitting wire using standard tri-linear interpolation and shell averaged tri-linear interpolation. A reference solution obtained by NEC-3 gives 3.53 MHz and 72.2 Ω

8.10.3 Loop antenna

In this section the thin-wire model is used to simulate a loop antenna in receiving mode. This experiment was also performed by Holland et al. [49] using two stair-cased wire approximations, one area matched, and one circumference matched. In both cases differences in the order of ten percent were reported due to staircasing effects.

The loop wire in [49] was discretized using 16 cells across the loop diameter and hence we use 50 segments for the loop antenna using the arbitrarily oriented wire model. The FD-TD grid contains $32 \times 32 \times 32$ cells with a cell size of 0.0625 m in cube. The loop diameter is one meter and the wire radius is 0.00779, 0.00286, and 0.00105 meter. This corresponds to Ω equal to 12, 14, and 16 where Ω is the fatness factor defined through

$$\Omega = 2 * \ln\left(\frac{\pi D}{a}\right), \quad (8.54)$$

where D is the loop diameter and a is the wire radius.

The loop antenna is oriented in the x-y plane and a plane wave propagating in the x-y plane impinges on the loop antenna with vertical magnetic polarization. The problem setup and the response after roughly 100 time steps is shown in Figure 8.27.




Figure 8.27. The loop antenna setup. An impinging plane wave excites the wire. The current is probed at the right-most segment, i.e. 135 degrees from the head-on segment. The cutting plane displays the magnitude of the electric field. The current is shown on the wire, whose radius is exaggerated in the figure.

The probed current is expanded in a Prony series and the mode corresponding to the 100 MHz mode in [49] is extracted. This mode corresponds to the one wavelength resonance and we can directly compare the results obtained by our shell interpolated wire model with results obtained using the Method of Moments (MoM). Holland et al. gives the frequency, damping coefficient, amplitude and phase for this case. But since their incident plane wave is not specified in space we cannot compare amplitude and phase. Also the damping is reported to be not fully

trustworthy and hence the MoM code developed within the GEMS project [39] is also utilized. In this MoM code we can choose between two wire models, a simple model and a refined model developed by Mazari. We also use two wire discretizations, one containing 50 segments and one built with 100 segments. The comparison between our new FD-TD wire model and the different MoM simulations is plotted in Figure 8.28.




Figure 8.28. Resonance frequencies for the new FD-TD model and five MoM results. The diamond results are found in [49] and are obtained by MoM. The other four types are different MoM solutions obtained by reading the zero crossing of the input impedance for a transmitting antenna. Ω is the fitness factor (8.54).

The variation in the MoM results indicates that results obtained by the shell interpolated wire model are accurate within one percent.

8.10.4 The wire length

A very interesting question arises when we compare the expected resonance frequencies in Table 8.2 with those obtained by Prony's method, i.e. Table 8.3. The differences are a few percent and it is tempting to believe that this is explained by the assumptions in the wire model. The effect seems to be that the modeled wire is longer than the discrete wire. In this case roughly a half segment longer at each wire end. However, in the last section of Chapter 4.6 in Balanis book on antenna theory [7] we find an interesting passage:

Depending on the radius of the wire, the length of the dipole for first resonance is about $l = 0.47 \lambda$ to 0.48λ ; the thinner the wire, the closer the length is to 0.48λ . Thus, for thicker wires, a larger segment of the wire has to be removed from $\lambda/2$ to achieve resonance.

Unfortunately, Balanis does not go further into this subject but leaves the reader to this uncertain range of a few percent mismatch in resonance frequency.

Another effect which might add to the understanding is based on the model for damped oscillations.

A free oscillation without damping can be modeled by the following differential equation:

$$\ddot{x} + \omega_0^2 x = 0, \quad (8.55)$$

which has the solution

$$x = C e^{i(\omega_0 t + \varphi)}. \quad (8.56)$$

Likewise a free oscillation with weak damping can be modeled by

$$\ddot{x} + 2\gamma\dot{x} + \omega_0 x = 0, \quad (8.57)$$

which has the solution

$$x = C e^{-\gamma t + i(\omega_d t)}, \quad (8.58)$$

where the eigenfrequency ω_d is given by

$$\omega_d = \sqrt{\omega_0^2 - \gamma^2}. \quad (8.59)$$

The resonant wavelength in Figure 8.17 and Table 8.2 are valid for undamped wire currents/voltages but for damped responses we have to compensate for the shift in the eigenfrequencies given by (8.59).

Using observed dampings for the receiving dipole antenna, we can determine the analytic damped frequency f_c according to (8.59). If we put the undamped analytic frequencies, the compensated analytic frequencies, and the observed frequencies together in a table, we can easily compare the effect of taking damping into account, see Table 8.6.

From Table 8.6 it is obvious that the damping does not explain the shift in frequency compared to expected frequencies. The compensation effect adjusts the undamped frequency less than one tenth of a percent (see Table 8.2) whereas the

mode	Analytic, undamped	Analytic, compensated for damping			Observed	
	f_0 [MHz]	γ_m [$\frac{1}{\mu s}$]	f_c [MHz]	$\frac{f_0-f_c}{f_0}$ %	f_m [MHz]	$\frac{f_0-f_m}{f_0}$ %
1	3.659	0.897	3.656	0.08	3.558	2.7
2	7.317	1.220	7.315	0.04	7.203	1.6
3	10.976	1.423	10.973	0.02	10.850	1.1
4	14.634	1.573	14.632	0.01	14.480	1.1
5	18.293	1.691	18.291	0.01	18.100	1.1
6	21.951	1.785	21.949	0.01	21.700	1.1
7	25.610	1.859	25.608	0.01	25.280	1.3
8	29.268	1.919	29.267	0.01	28.830	1.5

Table 8.6. The observed and corrected frequencies and wavelenghts due to the damping.

observed frequencies differ 1-3 percent from the analytic results. However, it is interesting to note that also the Method of Moments solution used as reference solution in Section 8.10.2 has a too low resonance frequency, 3.53 MHz for the first mode.

But if we take the frequency shift of 4 to 6 % into account, proposed by Balanis [7], our results are in good agreement with the “shifted” frequencies.

8.11 Conclusion

We have introduced a new and accurate subcell model to deal with arbitrarily oriented thin-wires. The model is useful for treating wires that cannot be aligned with the Cartesian grid, for example circular loop antennas and tilted straight wires.

Without the new model, staircasing is the traditional way to treat arbitrarily oriented thin-wires in FD-TD. The errors introduced by the staircasing can be several tens of percent. A regular tri-linear interpolation technique is investigated in order to allow arbitrarily oriented thin-wires. However, results obtained shows that the coupling between the Maxwell equations and the wire equations does not give consistent results for small changes in the wire position and orientation. These errors can also be several tens of percent.

For our new shell averaged tri-linear interpolation technique results are presented that illustrate the success of the new model. Both straight wires and circular loop antennas are used in the simulations. Induced currents and voltages, impedances and admittances, and resonance frequencies are confined within one percent deviation from each other. Our results are also in very good agreement with Method of Moments solutions.

Finally, a stability analysis for a periodic model problem is provided where the thin-wire model is shown to be stable under a time step limit.

Chapter 9

Color Electromagnetics

When a simulation is completed there is usually a need for a visualization of the result. In a few rare cases, only a number is needed from the computation, such as the monostatic radar cross section of a target at one frequency or the maximum field strength that appears in the simulation. But in most cases one has to comprehend a larger amount of data to draw meaningful conclusions from the computation. Hence, visualization of CEM computations is important. We can also couple the computation to a visualization system and thereby get the data on screen during simulation. This is only feasible for small problems, because the visualization engines are limited in performance compared to the rate of data output from the computation.

Remembering the target group is very important, because it affects the choice of visualization. If the target group are scientists, a regular x-y plot of derived quantities might reveal details of the method or simulation best; but other groups like CEM (*Color* ElectroMagnetics) more than regular x-y plots.

If you are restricted to 3D visualizations on a standard display there are some techniques that can be used to enhance the perception of 3D features. These include light sources, perspective rendering and making distant objects dimmer (called depth cue). These techniques are important because the human brain uses the effects to interpret the internal relations between geometrical objects.

A problem with visualizations on non-interactive media like paper is that interesting features may be concealed by objects or cutting planes and isosurfaces closer to the viewer. Semi-transparent color maps can decrease the hiding properties of these objects somewhat, and the viewer can see through them and thereby get a better experience of the data.

Another emerging technology is virtual reality, which provides tools and display systems for an immersive exploration of data. Used correctly, it really enhances the possibility of sharing results with others, especially if a complicated geometry is included. We have used the VR-Cube at PDC in this way, and also to understand the intricate interplay between the primary and dual grids used in the 3D Finite Volume solver described in Chapter 7.




Figure 9.1. An electromagnetic pulse strikes a metallic cavity with a narrow opening of $1/40$ of the diameter. Even though the opening is small a substantial amount of the pulse sneaks inside. The pictures shows the amplitude of the magnetic field at three different times. 4096×2048 quadratic cells are used (symmetry is utilized) which makes it possible to resolve the crack with 40×40 cells. This picture is a nice example of the perceptual enhancement of wave propagation by using a carpet plot. In this case the wave amplitude is mapped to both color and height.




Figure 9.2. This figure shows, for a distinct time, the solution of an open waveguide problem. An antenna is attached to an E-shaped waveguide and a current pulse is injected into the antenna. The antenna is modeled with a thin wire model. The calculation reveals how large fraction of energy that propagates along the waveguide compared to the energy that leaks out of the waveguide. The depth in the figure is enhanced by perspective visualization. Also note the depth cue, i.e. the colors fades away with depth in the image.




Figure 9.3. This picture shows the setup for a study of the current induction on a wire connecting two electronic devices in the presence of a transmitting cellular phone. The antenna on top of the telephone is driven by a current pulse and the picture shows a snapshot in time of the response. In this picture a semi transparent vertical cutting plane clearly illustrates the wave propagative nature of the transmitted waves. Opacity is lower for weaker field values, hence the objects behind the cutting plane is clearly visible.




Figure 9.4. The surface currents on the generic aircraft `Rund` is plotted a short time after a plane wave impinged the nose. The electric polarization of the plane wave is vertical (top) and horizontal (bottom). The 3D FD-FV hybrid solver described in Chapter 7 is used in this simulation and the hybrid mesh is shown in Figure 7.16. Plotting the surface current automatically reveals the PEC geometries. The geometric shape is easier to comprehend if a light model is used in the visualization. The light reflections in the fuselage gives a good perception of the curvature of the aircraft. Without shading (modeling directive light) the appearance would be very flat and the only visual clue to the shape would be our general experience of aircraft shapes.




Figure 9.5. This figure shows the surface currents on the SAAB 2000 aircraft 125 ns (1500 time steps) after a lightning stroke the nose. Also the magnitude of the magnetic field is shown on a cutting plane across the wings perpendicular to the fuselage. The maker of a picture must choose a good viewing direction, view angle and what part to visualize. These choices are often based on the desire to enhance interesting parts but sometimes also to suppress distracting features in the simulation. A minute spectator will notice that the outer boundary condition is working fine at the upper boundary but there is an anomaly at the lower boundary. A small mistake in the simulation setup gave a PEC condition at the lower boundary and hence the image was cropped to spare the spectator from trying to interpret this “anomaly” (read mistake).




Figure 9.6. The interior of the SAAB 2000 aircraft. Surface currents are shown at the same time as in the figure above. The view is from center of aircraft towards the cockpit. High surface currents are seen on the door pillar and the sill. To prevent cables from picking up high induced currents this type of simulations is important because proper countermeasures can be taken if the field distribution inside the aircraft is known. A sophisticated light model would take the blocking effect of the fuselage into account during rendering of the inside geometries. But this would slow down the rendering drastically and we have to cope with the results from the non-realistic light model used.




Figure 9.7. The 3D hybrid solver requires a hybrid grid. In this case an unstructured grid is used only in the vicinity of the nose of the SAAB 2000 aircraft and a structured grid is used for the rest of the problem. Different colors are assigned to different parts of the object. Part of the fuselage and unstructured grid is removed to let us see into the aircraft. Note the LEGO shape of the transition layer described in Section 7.5.




Figure 9.8. A thin wire is attached to the nose of the SAAB 2000 aircraft and a current injection is used to simulate a lightning strike. In this figure the surface current is plotted after a while. The electric field intensity is visualized on a cutting plane and also by using a volumetric rendering in the unstructured volume. The staggered components in FD-TD give rise to a half-cell bias between the field and geometry (see Section 4.8). This is manifested on the cutting plane as blue marks just above the fuselage.

Bibliography

- [1] S. Abarbanel and D. Gottlieb. A mathematical analysis of the PML method. *J. Comput. Phys.*, 134:357–363, 1997.
- [2] A. Akyurtlu, D. H. Werner, V. Veremey, D. J. Steich, and K. Aydin. Staircasing errors in FDTD at an air-dielectric interface. *IEEE Microwave Guided Wave Lett.*, 9(11):444–446, November 1999.
- [3] U. Andersson. Parallelization of a 3D FD-TD code for the Maxwell equations using MPI. In G. Kristensson, editor, *EMB 98 – Electromagnetic Computations for analysis and design of complex systems*, pages 94–101. SNRV, November 1998.
- [4] U. Andersson. Parallelization of a 3D FD-TD code for the Maxwell equations using MPI. In B. Kågström et al., editors, *Applied Parallel Computing, PARA'98*, Lecture Notes in Computer Science, No. 1541, pages 12–19, June 1998.
- [5] U. Andersson. *Time-Domain Methods for the Maxwell Equations*. TRITANA-0101, KTH, 2001.
- [6] U. Andersson, B. Engquist, G. Ledfelt, and O. Runborg. A contribution to wavelet-based subgrid modeling. *Applied and Computational Harmonic Analysis*, 7:151–164, 1999.
- [7] C. A. Balanis, editor. *Antenna theory: analysis and design*. John Wiley and Sons, Inc., USA, second edition, 1997.
- [8] J. T. Beale and A. Majda. Vortex methods II: Higher order accuracy in two and three dimensions. *Mathematics of Computation*, 39:29–52, 1995.
- [9] J.-P. Bérenger. A perfectly matched layer for the absorption of electromagnetic waves. *J. Comput. Phys.*, 114(1):185–200, 1994.
- [10] J.-P. Bérenger. A multiwire formalism for the FDTD method. *IEEE Trans. Electromagn. Compat.*, 42(3):257–264, August 2000.

- [11] M. Bingle. FDTD simulation of scattering by conducting wires in rectangular waveguide. In *Millennium Conference on Antennas and Propagation*, volume 1, page 92. European Space Agency, April 2000.
- [12] J. Bogerrod, A. Tijhuis, and J. Klaasen. Electromagnetic excitation of a thin wire: A traveling-wave approach. *IEEE Trans. Antennas Propagat.*, 46(8):1202–1211, August 1998.
- [13] V. A. Borovikov and B. Y. Kinber. *Geometrical Theory of Diffraction*. IEE, UK, 1994.
- [14] A. R. Bretones, R. Mittra, and R. G. Marti'n. A hybrid technique combining the method of moments in the time domain and FDTD. *IEEE Microwave Guided Wave Lett.*, 8(8):281–283, August 1998.
- [15] G. Burke and D. Steich. Extension of a thin-wire algorithm for wires moved laterally within a mesh. In *15th Annual Review of Progress in Applied Computational Electromagnetics*, volume 1, pages 404–411, Monterey, CA, March 1999.
- [16] CADfix. <http://www.cadfix.com/>.
- [17] A. C. Cangellaris and D. B. Wright. Analysis of the numerical error caused by the stair-stepped approximation of a conducting boundary in FDTD simulations of electromagnetic phenomena. *IEEE Trans. Antennas Propagat.*, 39(10):1518–1525, October 1991.
- [18] D. K. Cheng. *Field and Wave Electromagnetics*. Addison Wesley, second edition, 1989.
- [19] C. Christopoulos. *The Transmission-Line Modeling Method: TLM*. IEEE Press and Oxford University Press, 1995.
- [20] R. Coifman, V. Rokhlin, and S. Wandzura. The fast multipole method for the wave equation: A pedestrian approach. *IEEE Trans. Antennas Propagat.*, 35(7), 1993.
- [21] S. Dey and R. Mittra. A locally conformal finite-difference time-domain (FDTD) algorithm for modeling three-dimensional perfectly conducting objects. *IEEE Microwave Guided Wave Lett.*, 7(9):273–275, September 1997.
- [22] A. Ditkowski, K. Dridi, and J. S. Hesthaven. Convergent Cartesian grid methods for Maxwells equations in complex geometries. *J. Comput. Phys.*, 1999. Submitted.
- [23] S. Dodson, S. P. Walker, and M. J. Bluck. Implicitness and stability of time domain integral equation scattering analysis. *Appl. Comput. Electromagn. Soc. J.*, 13(3):291–301, May 1997.

- [24] A. Dominek, H. Shamanski, R. Wood, and R. Barger. A useful test body. In *Proc. Antenna Measurement Techniques Assoc.*, volume 24, September 1986.
- [25] A. K. Dominek and H. T. Shamanski. The almond test body. Technical Report 721929-9, The Ohio State University ElectroScience Laboratory, Department of Electrical Engineering, March 1990. prepared under Grant Number NSG1613, NASA Langley Research Center.
- [26] M. Douglas, M. Okoniewski, and M. A. Stuchly. Accurate modeling of thin-wire antennas in the FDTD method. *Microwave Opt. Technol. Lett.*, 21(4):261–265, May 1999.
- [27] K. H. Dridi, J. S. Hesthaven, and A. Ditkowski. Staircase free finite-difference time-domain formulation for general materials in complex geometries. *IEEE Trans. Antennas Propagat.*, 2001. Accepted.
- [28] F. Edelvik. Analysis of a finite volume solver for Maxwell’s equations. In R. Vilsmeier, editor, *Finite Volumes for Complex Applications II*, pages 141–148, Paris, France, July 1999. IVG, University of Duisburg, Hermes.
- [29] F. Edelvik. *Finite Volume Solvers for the Maxwell Equations in Time Domain*. Licentiate thesis No. 2000-005, Department of Information Technology, Uppsala University, October 2000.
- [30] F. Edelvik and B. Strand. Frequency dispersive materials for 3D hybrid solvers in time domain. In *Antenn 00, Nordic Antenna Symposium, Lund, Sweden*, pages 63–68, September 2000.
- [31] B. Engquist and A. Majda. Absorbing boundary conditions for the numerical simulation of waves. *Math. Comput.*, 31(139):629–651, 1977.
- [32] E. Engquist. Steering and visualization of electromagnetic simulations using globus. In B. Engquist, L. Johnsson, M. Hamill, and F. Short, editors, *Simulation and visualization on the grid*, Lecture Notes in Computational Science and Engineering. Paralleldatorcentrum, KTH, Springer, December 1999.
- [33] A. A. Ergin, B. Shanker, , and E. Michielssen. Time domain fast multipole methods: A pedestrian approach. *IEEE Antennas Propagat. Magazine*, 41(4):39–53, 1999.
- [34] A. A. Ergin, B. Shanker, and E. Michielssen. Fast evaluation of three-dimensional transient wave fields using diagonal translation operators. *J. Comput. Phys.*, 146(1):157–180, 1998.
- [35] Femlab, a Matlab toolbox. <http://www.comsol.se/>.
- [36] The Swedish Defence Research Establishment (FOA) which was superseded by the Swedish Defence Research Agency (FOI) january 2001. <http://www.foi.se>.

- [37] S. D. Gedney. Finite-difference time-domain analysis of microwave circuit devices on high performance vector/parallel computers. *IEEE Trans. Microwave Theory Tech.*, 43(10):2510–2514, October 1995.
- [38] S. D. Gedney. An anisotropic PML absorbing media for the FDTD simulation of fields in lossy and dispersive media. *Electromagnetics*, 16(4):399–415, 1996.
- [39] General Electromagnetic Solvers (GEMS). <http://www.psci.kth.se/Programs/GEMS/>.
- [40] M. Ghrist, B. Fornberg, and T. A. Driscoll. Staggered time integrators for wave equations. *SIAM J. Num. Anal.*, 38(3):718–741, 2000.
- [41] N. Gibbs, W. Poole, and P. Stockmeyer. An algorithm for reducing the bandwidth and profile of a sparse matrix. *SIAM J. Num. Anal.*, 13(2):236–250, April 1976.
- [42] M. J. Grote and J. B. Keller. Nonreflecting boundary conditions for Maxwell’s equations. *J. Comput. Phys.*, 139:327–342, 1998.
- [43] B. Gustafsson, H.-O. Kreiss, and J. Oliger. *Time Dependent Problems and Difference Methods*. Wiley-Interscience, 1995.
- [44] R. F. Harrington. *Field Computations by Moment Methods*. The Macmillan Co., New York, 1968.
- [45] L. He. FDTD - Advances in subsampling method, UPML and higher-order boundary conditions. Master’s thesis, University of Kentucky, 1997.
- [46] Encyclopedia Britannica Online, electromagnetic radiation. <http://www.eb.com:180/bol/topic?eu=108505&sctn=4>. Accessed January 9 2001.
- [47] High frequency structure simulator (HFSS). <http://www.ansoft.com>.
- [48] T. Hirono, Y. Shibata, W. W. Lui, S. Seki, and Y. Yoshikuni. The second-order condition for the dielectric interface orthogonal to the Yee-lattice axis in the FDTD scheme. *IEEE Microwave Guided Wave Lett.*, 10(9):359–361, September 2000.
- [49] R. Holland and L. Simpson. Finite-difference analysis of EMP coupling to thin struts and wires. *IEEE Trans. Electromagn. Compat.*, EMC-23(2):88–97, May 1981.
- [50] R. Holland and J. W. Williams. Total-field versus scattered-field finite-difference codes: A comparative assessment. *IEEE Trans. Nuclear Science*, NS-30(6):4583–4588, December 1983.

- [51] C.-T. Hwang and R.-B. Wu. Treating late-time instability of hybrid finite-element/finite-difference time-domain method. *IEEE Trans. Antennas Propagat.*, 47(2):227–232, February 1999.
- [52] T. Itoh and B. Houshmand, editors. *Time-Domain Methods for Microwave Structures: Analysis and Design*. IEEE Press, Piscataway, NJ, 1998.
- [53] J. B. Keller. Geometrical theory of diffraction. *J. opt. soc. of America.*, 52(2):61–72, February 1962.
- [54] D. Koh, H.-B. Lee, and T. Itoh. A hybrid full-wave analysis of via-hole grounds using finite-difference and finite-element time-domain methods. *IEEE Trans. Microwave Theory Tech.*, 45(12):2217–2223, 1997.
- [55] R. G. Kouyoumjian and P. H. Pathak. A uniform theory of diffraction for an edge in a perfectly conducting surface. *Proc. IEEE*, 62(11):1448–1461, 1974.
- [56] K. S. Kunz and R. J. Luebbers. *The Finite Difference Time Domain Method for Electromagnetics*. CRC Press, Boca Raton, FL, 1993.
- [57] Large Scale FD-TD, a PSCI project. <http://www.nada.kth.se/~ulfa/CEM.html>.
- [58] E. Larsson. *Domain Decomposition and Preconditioned Iterative Methods for the Helmholtz Equation*. PhD thesis, Uppsala University, 2000.
- [59] J.-F. Lee, R. Lee, and A. Cangellaris. Time-domain finite-element methods. *IEEE Trans. Antennas Propagat.*, 45(3):430–442, March 1997.
- [60] J.-F. Lee and Z. Sacks. Whitney elements time domain (WETD) methods. *IEEE Trans. Magnet.*, 31(3):1325–1329, 1995.
- [61] J.-Y Lee and N.-H. Myung. Locally tensor conformal FDTD method for modeling arbitrary dielectric surfaces. *Microwave Opt. Technol. Lett.*, 23(4):245–249, 1999.
- [62] J. E. Lumpp, Jr., S. K. Mazumdar, and S. D. Gedney. Performance modeling of the finite-difference time-domain method on parallel systems. *Applied Computational Electromagnetics Society Journal*, 13(2):147–159, 1998.
- [63] T. Martin. An improved near- to far-zone transformation for the finite-difference time-domain method. *IEEE Trans. Antennas Propagat.*, 46(9):1263–1271, September 1998.
- [64] T. Martin. *Broadband Electromagnetic Scattering and Shielding Analysis Using the Finite Difference Time Domain Method*. PhD thesis, Department of Physics and Measurement Technology, Linköpings universitet, Sweden, March 2001.

- [65] T. Martin and L. Pettersson. Dispersion compensation for Huygens' sources and far-zone transformation in FDTD. *IEEE Trans. Antennas Propagat.*, 48(4):494–501, April 2000.
- [66] J. C. Maxwell. *A Treatise on Electricity and Magnetism*. Clarendon Press, Oxford, 1873.
- [67] E. K. Miller. Pcs for ap and other em reflections. *IEEE Antennas Propagat. Magazine*, 41(3):83–88, June 1999.
- [68] P. Monk. A comparison of finite element methods for the time dependent Maxwell equations. *Mathematical and Numerical Aspects of Wave Propagation Phenomena*, G. Cohen et al. (eds), SIAM, pages 80–88, 1991.
- [69] P. Monk. An analysis of Nedelec's method for the spatial discretization of Maxwell's equations. *J. Comput. Appl. Math.*, 47:101–121, 1993.
- [70] P. Monk and K. Parrott. A dispersion analysis of finite element methods for Maxwell's equations. *SIAM J. Sci. Comput.*, 15(4):916–937, July 1994.
- [71] A. Monorchio and R. Mittra. A hybrid finite-element finite-difference time-domain (FE/FDTD) technique for solving complex electromagnetic problems. *IEEE Microwave Guided Wave Lett.*, 8(2):93–95, 1998.
- [72] T. G. Moore, J. G. Blaschak, A. Taflove, and G. A. Kriegsmann. Theory and application of radiation boundary operators. *IEEE Trans. Antennas Propagat.*, 36(12):1797–1812, Decemeber 1988.
- [73] The message passing interface (MPI) standard. <http://www-unix.mcs.anl.gov/mpi/>.
- [74] G. Mur. Absorbing boundary conditions for the finite-difference approximation of the time-domain electromagnetic-field equations. *IEEE Trans. Electromagn. Compat.*, EMC-23(4):377–382, November 1981.
- [75] NEC, Numerical Electromagnetic Code. <http://www.qsl.net/wb6tpu/swindex.html>.
- [76] J. C. Nedelec. Mixed finite elements in R3. *Num. Math.*, 35(9):315–341, September 1980.
- [77] C. Nordling and J. Österman. *Physics Handbook*. Studentlitteratur, 4th edition, 1987.
- [78] Swedish National Bard for Industrial and Technical Development (NUTEK) which was superseded by The Swedish Agency for Innovation Systems (VINNOVA) january 2001. <http://www.vinnova.se>.

- [79] M. Okoniewski, E. Okoniewska, and M. A. Stuchly. Three-dimensional sub-gridding algorithm for FDTD. *IEEE Trans. Antennas Propagat.*, 45(3):422–429, March 1997.
- [80] P. G. Petropoulos. Analysis of exponential time-differencing for FDTD in lossy dielectrics. *IEEE Trans. Antennas Propagat.*, 45(6):1054–1057, June 1997.
- [81] Parallel and Scientific Computing Institute (PSCI). <http://www.psci.kth.se/>.
- [82] O. M. Ramahi. The concurrent complementary operators method for FDTD mesh truncation. *IEEE Trans. Antennas Propagat.*, 46(10):1475–1482, October 1998.
- [83] D. J. Riley. Transient finite-elements for computational electromagnetics: Hybridization with finite differences, modeling thin wires and thin slots, and parallel processing. In *17th Annual Review of Progress in Applied Computational Electromagnetics*, Monterey, CA, March 2001.
- [84] D. J. Riley and C. D. Turner. VOLMAX: A solid-model-based, transient volumetric Maxwell solver using hybrid grids. *IEEE Antennas Propagat. Magazine*, 39(1):20–33, 1997.
- [85] D. J. Riley and C. D. Turner. The VOLMAX transient electromagnetic modeling system, including sub-cell slots and wires on random non-orthogonal cells. In *14th Annual Review of Progress in Applied Computational Electromagnetics*, volume 2, pages 816–824, Monterey, CA, March 1998.
- [86] A. Ruddle, D. Ward, R. Scarmuzza, and V. Trenkic. Development of thin wire models in tlm. In *IEEE Int. EMC Symp.*, volume 2, pages 196–201, Denver, CO, 1998.
- [87] T. Rylander. *The Application of Edge Elements in Electromagnetics*. Licentiate thesis, Chalmers University of Technology, 2000.
- [88] T. Rylander and A. Bondeson. Stable FEM-FDTD hybrid method for Maxwell’s equations. *Comput. Phys. Comm.*, 125:75–82, March 2000.
- [89] B. P. Rynne and P. D. Smith. Stability of time marching algorithms for the electric field integral equations. *J. Electromagn. Waves Applicat.*, 12:1181–1205, 1990.
- [90] S. Selleri, J. Y. Dauvignac, G. Pelosi, and C. Pichot. Comparison between FDTD and hybrid FDTD-FETD applied to scattering and antenna problems. *Microwave Opt. Technol. Lett.*, 18(4):247–250, 1998.

- [91] V. Shankar, W. Hall, and H. Mohammadian. A CFD-based finite-volume procedure for computational electromagnetics - interdisciplinary applications of CFD methods. *AIAA*, pages 551–564, 1989.
- [92] V. Shankar, W. F. Hall, and A. Mohammadian. A time-domain differential solver for electromagnetic scattering problems. *Proc. IEEE*, 77(5):709–721, 1989.
- [93] B. Shanker, A. A. Ergin, K. Aygun, and E. Michielssen. Analysis of transient electromagnetic scattering phenomena using a two-level plane wave time-domain algorithm. *IEEE Trans. Antennas Propagat.*, 48(4), April 2000.
- [94] P. P. Silvester and R. L. Ferrari. *Finite Elements for Electrical Engineers 2nd Ed.* Cambridge University Press, Cambridge, 1990.
- [95] J. M. Song and W. C. Chew. Multilevel fast-multipole algorithm for solving combined field integral equations of electromagnetic scattering. *Microwave Opt. Technol. Lett.*, 10(14), 1995.
- [96] M. W. Steeds. A comparison of two conformal methods for FDTD modeling. Master's thesis, Washington State University, Pullman, WA, 1995.
- [97] D. M. Sullivan. *Electromagnetic Simulation Using the FDTD Method.* IEEE Press, New York, 2000.
- [98] A. Taflove. *Computational Electrodynamics: The Finite-Difference Time-Domain Method.* Artech House, Boston, MA, first edition, 1995.
- [99] A. Taflove, editor. *Advances in Computational Electrodynamics: The Finite-Difference Time-Domain Method.* Artech House, Boston, MA, 1998.
- [100] A. Taflove. *Computational Electrodynamics: The Finite-Difference Time-Domain Method.* Artech House, Boston, MA, second edition, 2000.
- [101] C. H. Teng, A. Ditkowski, and J. S. Hesthaven. Modeling dielectric interfaces in the FDTD-method: A comparative study. *IEEE Trans. Antennas Propagat.*, 2000. Submitted.
- [102] A.-K. Tornberg. *Interface Tracking Methods with Application to Multiphase Flows.* TRITA-NA-0010, KTH, 2000.
- [103] K. R. Umashankar, A. Taflove, and B. Beker. Calculation and experimental validation of induced currents on coupled wires in an arbitrary shaped cavity. *IEEE Trans. Antennas Propagat.*, 35(11):1248–1257, November 1987.
- [104] J. L. Volakis, A. Chatterjee, and L. C. Kempel. *Finite Element Method for Electromagnetics, Antennas, Microwave Circuits and Scattering Applications.* IEEE Press, 1998.

- [105] J. J. H. Wang. *Generalized Moment Methods in Electromagnetics, Formulation and Computer Solution of Integral Equations*. John Wiley & Sons Inc, 1991.
- [106] S. Watanabe and M. Taki. An improved FDTD model for the feeding gap of a thin-wire antenna. *IEEE Microwave Guided Wave Lett.*, 8(4):152–154, April 1998.
- [107] A. J. Wlodarczyk, V. Trenkci, R. A. Scarmuzza, and C. Cristopoulos. A fully integrated multiconductor model for tlm. *IEEE Trans. Microwave Theory Tech.*, 46(12):2431–2437, December 1998.
- [108] R.-B. Wu and T. Itoh. Hybridizing FD-TD analysis with unconditionally stable FEM for objects of curved boundary. In *IEEE MTT-S Int. Microwave Symp.*, volume 2, pages 833–836, Orlando, FL, 1995.
- [109] R.-B. Wu and T. Itoh. Hybrid finite-difference time-domain modeling of curved surfaces using tetrahedral edge elements. *IEEE Trans. Antennas Propagat.*, 45(8):1302–1309, 1997.
- [110] K. S. Yee. Numerical solution of initial boundary value problems involving Maxwell’s equations in isotropic media. *IEEE Trans. Antennas Propagat.*, 14(3):302–307, March 1966.
- [111] K. S. Yee, J. S. Chen, and A. H. Chang. Conformal finite-difference time-domain (FDTD) with overlapping grids. *IEEE Trans. Antennas Propagat.*, 40(9):1068–1075, September 1992.
- [112] M. S. Yeung. Application of the hybrid FDTD-FETD method to dispersive materials. *Microwave Opt. Technol. Lett.*, 23(4):238–242, 1999.