MASTER'S THESIS

High-Performance Finite Element Simulation of Electromagnetic Waves for Ion Cyclotron Resonance Heating

Adam Canarutto

M2 Physique Fondamentale – Plasma and Fusion
Aix-Marseille Université (AMU)

Host Institution

Barcelona Supercomputing Center (BSC-CNS)
Fusion Group

Supervisors

Hernán Domingo Ramos
Alejandro Soba
Ezequiel Goldberg

Academic Year 2025–2026

Abstract

Ion Cyclotron Resonance Heating (ICRH) is one of the main auxiliary heating methods employed in magnetic-confinement fusion devices. It is based on the interaction of radio-frequency electromagnetic waves with magnetized plasmas, involving phenomena such as wave–particle interactions, anisotropic material response, and spatially varying plasma properties. The numerical modelling of these systems requires methods capable of describing electromagnetic wave propagation in complex geometries and, for large-scale configurations, can require substantial computational resources. To study numerically those complex effects, the Fusion Group of the Barcelona Supercomputing Center has created `EMWAVE`, a full-wave electromagnetic solver developed within Alya, a multiphysics finite-element framework also developed at the Barcelona Supercomputing Center. Alya is designed for massively parallel simulations on high-performance computing architectures, providing the computational framework required for large-scale electromagnetic problems relevant to fusion applications.

This thesis focuses on the extension, verification, and computational assessment of the `EMWAVE` solver. Its numerical capabilities were extended through the implementation and assessment of quadrilateral finite elements and second-order triangular elements. The electromagnetic material model was extended to account for electrical conductivity through an effective complex permittivity formulation. In addition, the post-processing capabilities of the solver were expanded beyond the electric field by implementing the computation of the magnetic flux density $\mathbf{B}$, the magnetic field $\mathbf{H}$, and the current density $\mathbf{J}$.

The implemented developments were verified and validated through dedicated benchmark cases and quantitative comparisons with the ERMES 20.0 electromagnetic solver. The computational behaviour of `EMWAVE` was further investigated through strong-scaling analyses performed on the MareNostrum 5 supercomputer. Finally, the solver was coupled with the DAKOTA toolkit to perform sensitivity analyses of electromagnetic simulations. The sensitivity study also investigated the influence of spatial resolution, showing the importance of using a sufficiently resolved discretization before interpreting the response of the electromagnetic solution to variations of physical input parameters.

The developments and analyses presented in this thesis provide a more complete and thoroughly assessed version of `EMWAVE`, supporting its continued development for the simulation of electromagnetic wave propagation in fusion plasmas.

Contents

1 Introduction **4**
1.1 Nuclear Fusion as a Sustainable Energy Source . . . 4
1.2 Plasma Heating in Magnetic Confinement Devices . . . 4
1.2.1 Radio-Frequency Heating . . . 4
1.3 Objectives of the Thesis . . . 5

2 Theoretical Background **7**
2.1 Maxwell's Equations . . . 7
2.2 Time-Harmonic Formulation . . . 7
2.3 Wave Equation . . . 8
2.3.1 General Scalar Form . . . 9
2.4 Cold Plasma Approximation . . . 10
2.5 Conductive Media . . . 12
2.6 Wave Propagation in Magnetized Plasmas . . . 13
2.6.1 Ordinary Wave . . . 14
2.6.2 Extraordinary Wave . . . 15

3 Finite Element Method **17**
3.1 Theoretical Foundations of FEM . . . 17
3.2 Galerkin Finite Element Formulation . . . 17
3.3 Element Categories . . . 18
3.3.1 Linear Triangular Element . . . 20
3.3.2 Quadratic Triangular Element . . . 21
3.3.3 Bilinear Quadrilateral Element . . . 22
3.4 Weak Formulation . . . 23
3.5 Edge Elements . . . 24
3.5.1 Nédélec Basis Functions . . . 25
3.5.2 Weak Formulation for Edge Elements . . . 26

4 Developments in the EMWAVE Module **28**
4.1 Alya and the EMWAVE Module . . . 28
4.2 Input Files and Problem Definition . . . 28
4.3 Computational Workflow of EMWAVE . . . 28
4.4 Extension to Quadratic and Quadrilateral Elements . . . 29
4.5 Reference Cold-Plasma Case . . . 30
4.6 Quadratic Triangular Elements . . . 31
4.7 Bilinear Quadrilateral Elements . . . 32
4.8 Comparison of the Discretizations . . . 33
4.9 Implementation of Electrical Conductivity . . . 35
4.10 Implementation of Additional Electromagnetic Fields . . . 36

5 Verification and Validation **39**
5.1 Conductive Rectangular Waveguide . . . 39
5.2 Numerical Comparison against ERMES . . . 40

6 Scalability Analysis on MareNostrum 5 **46**
6.1 Strong Scaling Methodology . . . 46
6.2 Scaling Results . . . 47
6.2.1 `TS-CIRC_DIELECTRIC_ISOTROPIC` . . . 47
6.2.2 `TS-CUBE_COLDPLASMA_ABC` . . . 48
6.2.3 Discussion . . . 49

7 Sensitivity Analysis with DAKOTA **51**
7.1 Sensitivity Analysis and the DAKOTA Toolkit 51
7.2 Coupling DAKOTA with Alya/EMWAVE 51
7.3 Reference Electromagnetic Problem 52
7.4 Initial Sensitivity Analysis on the Original Mesh 52
7.5 Sensitivity Analysis on the Refined Mesh 53

8 Conclusions **56**

1 Introduction

1.1 Nuclear Fusion as a Sustainable Energy Source

The continuous growth of global energy demand, together with the urgent need to reduce greenhouse gas emissions, has intensified the search for sustainable and low-carbon energy sources [1]. While renewable technologies such as solar and wind power are expected to play a major role in the future energy mix, their intermittent nature poses significant challenges for the stability of electrical grids. Consequently, considerable research efforts are being devoted to the development of complementary energy sources capable of providing large amounts of reliable and carbon-free electricity.

Among the various alternatives currently under investigation, controlled thermonuclear fusion is widely regarded as one of the most promising long-term solutions. Fusion is the process that powers the Sun and other stars, where light atomic nuclei combine to form heavier nuclei, releasing large amounts of energy as a consequence of the mass difference between reactants and products. Unlike nuclear fission, fusion does not rely on self-sustaining chain reactions and offers several advantages, including the potential for a sustainable fuel cycle based on deuterium extracted from seawater and tritium bred from lithium, intrinsically safe operation, and a significantly lower production of long-lived radioactive waste [4].

Among the various fusion reactions under investigation, the deuterium–tritium (D–T) reaction is currently considered the most suitable for first-generation fusion power plants because it exhibits the highest fusion cross-section at the lowest achievable plasma temperatures [5]. The reaction can be written as

$$^2\mathrm{H} + {}^3\mathrm{H} \rightarrow {}^4\mathrm{He} + n + 17.6\ \mathrm{MeV}, \tag{1.1}$$

where approximately 80% of the released energy is carried away by the neutron, while the remaining 20% is deposited in the alpha particle (^{4}He). However, overcoming the electrostatic repulsion between positively charged nuclei requires plasma temperatures of the order of 10^8 K, with no known material capable of withstanding such temperatures through direct contact.

To achieve such extreme conditions, several confinement concepts have been proposed. Among them, *magnetic confinement fusion* has emerged as one of the most mature and extensively investigated approaches. In magnetic confinement devices, strong magnetic fields are used to confine the charged particles, preventing direct contact between the plasma and the reactor walls while maintaining the high temperatures required for fusion reactions.

The most widely studied magnetic confinement configuration is the tokamak, although alternative concepts such as stellarators have also demonstrated significant potential in recent years [3]. Regardless of the specific configuration, all magnetic confinement devices face the common challenge of efficiently heating the plasma to fusion-relevant temperatures while maintaining plasma stability and minimizing energy losses.

1.2 Plasma Heating in Magnetic Confinement Devices

During the plasma start-up phase, in which the neutral fuel gas gets fully ionized, the toroidal plasma current generates heat through Joule dissipation, with a volumetric power density given by

$$P_{\mathrm{Ohmic}} = \mathbf{J} \cdot \mathbf{E} = \eta J^2, \tag{1.2}$$

where $\mathbf{J}$ is the plasma current density, $\mathbf{E}$ is the electric field, and η is the plasma resistivity. However, the plasma resistivity decreases approximately according to the Spitzer law ($\eta \propto T^{-3/2}$) as the temperature increases [6]. Consequently, the Ohmic heating power rapidly decreases and becomes insufficient to raise the plasma temperature to the values required for efficient fusion reactions.

To overcome this limitation, modern magnetic confinement devices employ several auxiliary heating techniques. The most widely adopted methods are Neutral Beam Injection (NBI), Electron Cyclotron Resonance Heating (ECRH), and Ion Cyclotron Resonance Heating (ICRH). While NBI transfers energy by injecting energetic neutral particles into the plasma, both ECRH and ICRH belong to the family of radio-frequency (RF) heating systems, in which electromagnetic waves are launched into the plasma and transfer energy through resonant wave–particle interactions.

1.2.1 Radio-Frequency Heating

The operating principle of radio-frequency heating is based on the interaction between electromagnetic waves and charged particles gyrating around magnetic field lines. The angular cyclotron frequency of a particle species s is given by

$$\omega_{c,s} = \frac{|q_s|B}{m_s}, \tag{1.3}$$

where q_s and m_s are the charge and mass of the particle species, respectively, and B is the magnitude of the magnetic flux density. Since the ion mass is typically three to four orders of magnitude larger than the electron mass (for example, the proton-to-electron mass ratio is approximately 1836), the ion cyclotron frequency is correspondingly much lower than the electron cyclotron frequency.

This difference naturally leads to two complementary radio-frequency heating techniques. Electron Cyclotron Resonance Heating (ECRH) operates at frequencies matching the electron cyclotron frequency, typically in the microwave range of tens to hundreds of GHz. The injected electromagnetic waves are primarily absorbed by the electrons, which subsequently transfer part of their energy to the ions through Coulomb collisions.

In contrast, Ion Cyclotron Resonance Heating (ICRH) operates at frequencies close to the ion cyclotron frequency, generally in the range of several tens of MHz. In this case, electromagnetic energy can be transferred efficiently to the ion population through resonant interactions at the ion cyclotron frequency and its harmonics [7]. Both RF heating systems, together with the Neutral Beam Injection system introduced previously, are illustrated for the ITER configuration in Figure 1. Although both ECRH and ICRH are well-established and complementary plasma heating techniques, they address different operational requirements. ECRH is particularly suited for localized heating and current profile control, whereas ICRH provides efficient ion heating and is widely employed to increase the thermal energy of the plasma [7].

The numerical modelling of RF wave propagation is therefore an important component in the study and design of plasma heating systems. In realistic fusion devices, the interaction between the launched electromagnetic waves, the plasma, and the surrounding structures is influenced by the complex device geometry, the magnetic-field configuration, the plasma properties, and the electromagnetic properties of the materials. These features make analytical descriptions insufficient for many realistic configurations and motivate the use of numerical methods capable of representing complex geometries and heterogeneous or anisotropic media.

The Finite Element Method (FEM) is particularly well suited to this class of electromagnetic problems because of its flexibility in the treatment of complex geometries, material properties, and boundary conditions [8]. However, the spatial resolution required to represent electromagnetic wave propagation in large and geometrically complex fusion configurations can lead to numerical problems involving very large numbers of degrees of freedom. Efficient parallel algorithms and high-performance computing resources therefore become important for performing such simulations at the required spatial resolution.

For this purpose, the Fusion Group at the Barcelona Supercomputing Center has developed `EMWAVE` [9], an electromagnetic module integrated within the Alya multiphysics simulation framework [19]. The module provides a finite-element framework for the simulation of radio-frequency electromagnetic wave propagation in magnetically confined fusion plasmas and is described in detail in Chapter 4. The present work uses `EMWAVE` to extend and assess electromagnetic modelling capabilities relevant to ICRH applications, providing the motivation for the numerical developments presented throughout this thesis.

1.3 Objectives of the Thesis

This Master's thesis was carried out within the Fusion Group of the Barcelona Supercomputing Center (BSC-CNS), the Spanish National Supercomputing Center and a research institution specializing in high-performance computing and computational science. The Fusion Group is part of the Computer Applications in Science and Engineering (CASE) Department and focuses on the development and application of numerical methods and high-performance computing tools for the modelling of magnetic-confinement fusion devices.

The main contributions for the `EMWAVE` module of this work can be summarized as follows:

- the implementation and assessment of quadrilateral finite elements and second-order triangular elements (see section 4.4);
- the extension of the electromagnetic material model to account for electrical conductivity through a complex permittivity formulation (see section 4.9);
- the implementation of additional electromagnetic post-processing quantities, namely the magnetic flux density **B**, the magnetic field strength **H**, and the current density **J** (see section 4.10);
- the verification and validation of the implemented developments through benchmark test cases and comparisons with reference solutions and the ERMES [20] electromagnetic solver (see Chapter 5);

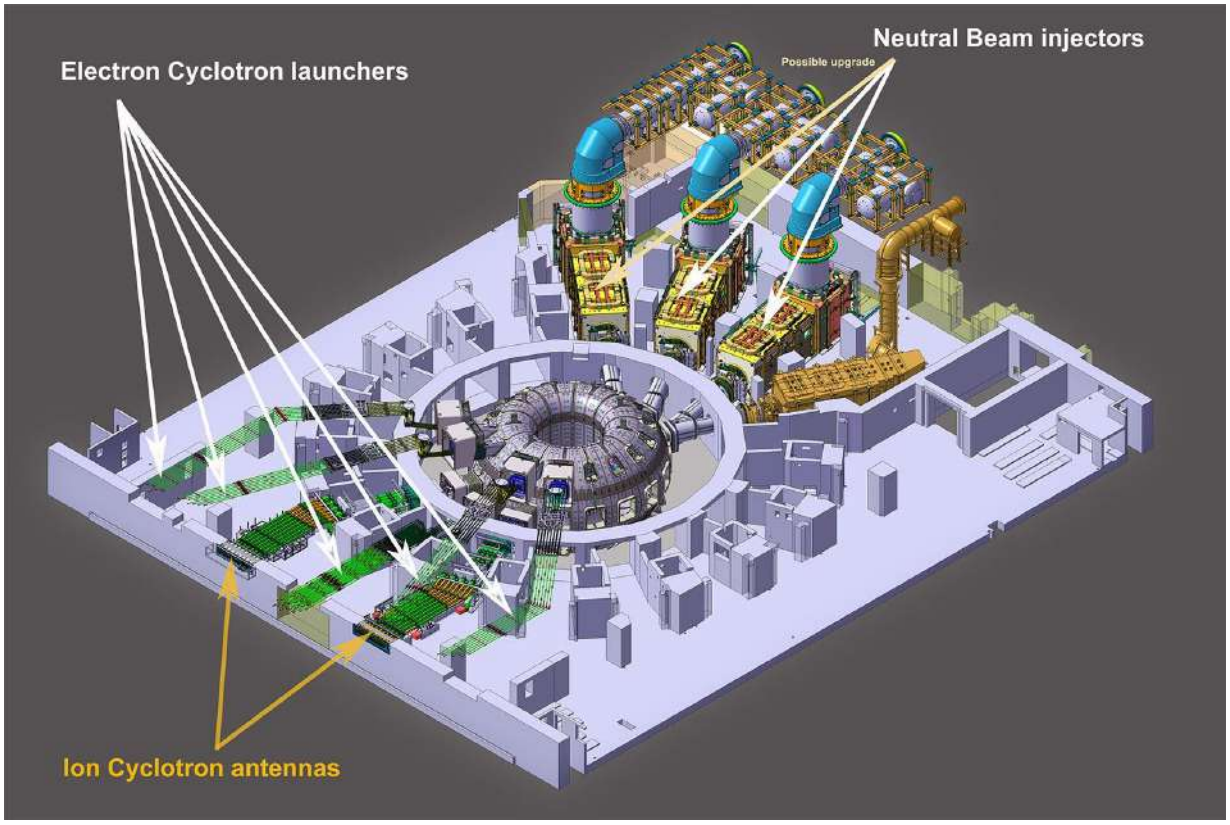

Figure 1: Main external heating systems of ITER: Electron Cyclotron Resonance Heating (ECRH), Neutral Beam Injection (NBI), and Ion Cyclotron Resonance Heating (ICRH). Source: ITER Organization [2].

- the assessment of the parallel performance of `EMWAVE` on the MareNostrum 5 supercomputer through strong-scaling analyses (see Chapter 6);
- the integration of the Dakota toolkit [13] to perform sensitivity analyses of electromagnetic simulations (see Chapter 7).

Together, these developments aim to extend the numerical capabilities of `EMWAVE`, verify the correctness of the newly implemented features, and assess their behaviour in computationally demanding electromagnetic simulations.

The remainder of this thesis is organized as follows. Chapter 2 introduces the theoretical background required for the electromagnetic and numerical formulations considered in this work, including electromagnetic wave propagation in magnetized plasmas. Chapter 3 presents the theoretical foundations of the Finite Element Method, including nodal and edge element formulations. Chapter 4 describes the Alya multiphysics framework and the `EMWAVE` electromagnetic module, and presents the numerical and physical developments implemented in `EMWAVE`, while Chapter 5 is devoted to their verification and validation against ERMES 20.0. Chapter 6 investigates the parallel performance of the solver on MareNostrum 5, and Chapter 7 presents the coupling with Dakota and the resulting sensitivity analysis. Finally, Chapter 8 summarizes the main conclusions of the work and discusses possible directions for future development.

2 Theoretical Background

This chapter introduces the theoretical concepts required to understand the numerical developments presented in this work. The governing equations describing electromagnetic wave propagation are first introduced, followed by the assumptions adopted for modelling magnetized fusion plasmas. The resulting cold-plasma dielectric response and its influence on wave propagation are then discussed, providing the physical framework for the numerical developments presented in the following chapters.

2.1 Maxwell's Equations

The propagation of electromagnetic waves in fusion plasmas is governed by Maxwell's equations, which provide a complete macroscopic description of the interaction between electric and magnetic fields, electric charges, and electric currents. Since the objective of this work is the numerical simulation of radio-frequency electromagnetic waves generated by ICRH antennas, Maxwell's equations constitute the fundamental mathematical model solved by the `EMWAVE` module.

In differential form, Maxwell's equations can be written as

$$\nabla \cdot \mathbf{D} = \rho, \tag{2.1}$$
$$\nabla \cdot \mathbf{B} = 0, \tag{2.2}$$
$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}, \tag{2.3}$$
$$\nabla \times \mathbf{H} = \mathbf{J} + \frac{\partial \mathbf{D}}{\partial t}, \tag{2.4}$$

where $\mathbf{E}$ and $\mathbf{H}$ denote the electric and magnetic field intensities, $\mathbf{D}$ and $\mathbf{B}$ are the electric displacement and magnetic flux density, respectively, ρ is the electric charge density, and $\mathbf{J}$ is the current density.

The electromagnetic fields are related through the constitutive equations

$$\mathbf{D} = \varepsilon_0 \varepsilon_r \mathbf{E}, \tag{2.5}$$
$$\mathbf{B} = \mu_0 \mu_r \mathbf{H}, \tag{2.6}$$

where ε_0 and μ_0 are the vacuum permittivity and permeability, while ε_r and μ_r denote the relative permittivity and permeability of the medium [14].

In the case of magnetized plasmas, however, the electromagnetic response is generally anisotropic and dispersive. As a consequence, the scalar relative permittivity must be replaced by a dielectric tensor whose components depend on the plasma properties, the applied magnetic field, and the wave frequency. This formulation will be introduced in section 2.3.

2.2 Time-Harmonic Formulation

In the present work, the antenna excitations are assumed to be monochromatic, with an angular frequency ω. This approximation allows to treat Maxwell's equations in the frequency domain rather in the time domain. Under this assumption, any electromagnetic quantity $f(\mathbf{x}, t)$ can be expressed as

$$f(\mathbf{x}, t) = \Re\left\{\hat{f}(\mathbf{x}) e^{-i\omega t}\right\}, \tag{2.7}$$

where $\hat{f}(\mathbf{x})$ denotes the complex amplitude (or phasor) of the field, and $\Re\{\cdot\}$ represents the real part.

Using the standard phasor representation of time-harmonic electromagnetic fields [14], the time derivative is replaced by

$$\frac{\partial}{\partial t} \longrightarrow -i\omega, \tag{2.8}$$

yielding the frequency-domain form of Maxwell's equations,

$$\nabla \times \mathbf{E} = i\omega \mathbf{B}, \tag{2.9}$$
$$\nabla \times \mathbf{H} = \mathbf{J} - i\omega \mathbf{D}, \tag{2.10}$$
$$\nabla \cdot \mathbf{D} = \rho, \tag{2.11}$$
$$\nabla \cdot \mathbf{B} = 0. \tag{2.12}$$

The use of complex phasors is standard in computational electromagnetics, as it transforms a time-dependent problem into a stationary boundary-value problem. Consequently, for a linear system excited at a prescribed frequency, only one complex-valued system must be solved to determine the steady-state electromagnetic response. This generally requires a lower computational cost than a time-domain simulation, which must instead resolve the temporal evolution of the fields over multiple oscillation periods. The frequency-domain approach is therefore particularly suitable for the present work, in which the antenna operates at a fixed frequency and transient phenomena are not considered.

2.3 Wave Equation

Starting from the frequency-domain Maxwell's equations, it is possible to derive a single vector equation for the electric field by eliminating the magnetic field. This equation constitutes the starting point for many finite-element formulations employed in computational electromagnetics [8].

Starting from the differential Faraday equation,

$$\nabla \times \mathbf{E} = i\omega \mathbf{B} = i\omega \mu \mathbf{H}, \tag{2.13}$$

it follows that,

$$\mathbf{H} = \frac{1}{i\omega\mu} \nabla \times \mathbf{E}. \tag{2.14}$$

Substituting this expression into eq. (2.10) yields

$$\nabla \times \left(\frac{1}{\mu} \nabla \times \mathbf{E} \right) + i\omega \mathbf{J} - \omega^2 \varepsilon \mathbf{E} = 0 \tag{2.15}$$

Using the constitutive relations

$$\mu = \mu_0 \mu_r, \qquad \varepsilon = \varepsilon_0 \varepsilon_r, \tag{2.16}$$

where μ_0 and ε_0 denote the permeability and permittivity of free space, respectively, while μ_r and ε_r are the corresponding relative material properties, and multiplying the equation by μ_0, the wave equation becomes

$$\nabla \times \left(\mu_r^{-1} \nabla \times \mathbf{E} \right) - \omega^2 \mu_0 \varepsilon_0 \varepsilon_r \mathbf{E} = -i\omega \mu_0 \mathbf{J}. \tag{2.17}$$

Introducing the free-space wavenumber,

$$k_0 = \omega \sqrt{\mu_0 \varepsilon_0}, \tag{2.18}$$

and the intrinsic impedance of free space,

$$\eta_0 = \sqrt{\frac{\mu_0}{\varepsilon_0}}, \tag{2.19}$$

the vector wave equation can be written in the compact form

$$\nabla \times \left(\mu_r^{-1} \nabla \times \mathbf{E} \right) - k_0^2 \varepsilon_r \mathbf{E} = -ik_0 \eta_0 \mathbf{J}. \tag{2.20}$$

One can also show that, starting from the same eq. (2.10), but isolating $\mathbf{E}$ instead of $\mathbf{H}$, the following equation can be found:

$$\nabla \times \left(\varepsilon^{-1} \nabla \times \mathbf{H} \right) - \omega^2 \mu \mathbf{H} = \nabla \times \left(\varepsilon^{-1} \mathbf{J} \right) \tag{2.21}$$

Equation (2.20) and eq. (2.21) are both 3D vector wave equations. However, for an invariant geometry along the z direction, and assuming that the current density $\mathbf{J}$ and the material properties do not depend on z, the three-dimensional electromagnetic problem can be reduced to a two-dimensional formulation [8].

It is also possible to perform a decomposition in the two fundamental polarizations: TE and TM [8]. These two acronyms stand for *transverse electric* and *transverse magnetic* respectively, and correspond to the cases $E_z = 0, H_z \neq 0$, for TE, and $E_z \neq 0, H_z = 0$ for TM. By doing so, we retrieve the *Helmholtz equations*, that read:

$$\nabla \cdot (\Lambda_\mu \cdot \nabla E_z) \; + \; k_0^2 \varepsilon_{zz} E_z = 0 \tag{2.22}$$

$$\nabla \cdot (\Lambda_\varepsilon \cdot \nabla H_z) \; + \; k_0^2 \mu_{zz} H_z = 0 \tag{2.23}$$

with,

$$\Lambda_\mu = \frac{(\boldsymbol{\mu}_r)^T_{sub}}{|(\boldsymbol{\mu}_r)^T_{sub}|} \text{ and } \Lambda_\varepsilon = \frac{(\boldsymbol{\varepsilon}_{rc})^T_{sub}}{|(\boldsymbol{\varepsilon}_{rc})^T_{sub}|},$$

where the relative permittivity and permeability tensors are written as

$$\boldsymbol{\varepsilon}_{rc} = \begin{pmatrix} \varepsilon_{xx} & \varepsilon_{xy} & 0 \\ \varepsilon_{yx} & \varepsilon_{yy} & 0 \\ 0 & 0 & \varepsilon_{zz} \end{pmatrix} = \begin{pmatrix} (\boldsymbol{\varepsilon}_{rc})_{\text{sub}} & \mathbf{0} \\ \mathbf{0}^T & \varepsilon_{zz} \end{pmatrix}, \tag{2.24a}$$

$$\boldsymbol{\mu}_r = \begin{pmatrix} \mu_{xx} & \mu_{xy} & 0 \\ \mu_{yx} & \mu_{yy} & 0 \\ 0 & 0 & \mu_{zz} \end{pmatrix} = \begin{pmatrix} (\boldsymbol{\mu}_r)_{\text{sub}} & \mathbf{0} \\ \mathbf{0}^T & \mu_{zz} \end{pmatrix}. \tag{2.24b}$$

Equation (2.20) and eq. (2.21) represent the general vector wave equations in the frequency domain for linear media. For isotropic materials, $\boldsymbol{\varepsilon_r}$ and $\boldsymbol{\mu_r}$ reduce to scalar quantities, whereas in anisotropic media they must be represented by tensors. In magnetized plasmas, however, the dielectric response is anisotropic and cannot be described by a scalar relative permittivity. Section 2.4 introduces the cold-plasma dielectric tensor, which extends this formulation to the case of magnetized fusion plasmas.

2.3.1 General Scalar Form

The TE and TM Helmholtz equations presented above can both be written in the following general form,

$$-\nabla \cdot (\mathbf{p} \cdot \nabla u) + qu = f, \tag{2.25}$$

where u denotes the unknown scalar field, $\mathbf{p}$ is a (possibly tensor-valued) material coefficient, q is a reaction coefficient, and f represents the source term.

For the TM polarization,

$$u = E_z, \tag{2.26}$$

$$\mathbf{p} = \boldsymbol{\Lambda}_\mu, \tag{2.27}$$

$$q = -k_0^2 \varepsilon_{zz}, \tag{2.28}$$

whereas, for the TE polarization,

$$u = H_z, \tag{2.29}$$

$$\mathbf{p} = \boldsymbol{\Lambda}_\varepsilon, \tag{2.30}$$

$$q = -k_0^2 \mu_{zz}. \tag{2.31}$$

The source term f depends on the particular excitation adopted in the electromagnetic problem. In the case of antenna or waveguide excitation, f accounts for the imposed source, whereas for plane-wave scattering problems it also contains the contribution associated with the known incident field. The compact notation introduced in eq. (2.25) considerably simplifies the derivation of the finite-element weak formulation presented in section 3.4.

2.4 Cold Plasma Approximation

Unlike conventional dielectric materials, a plasma is composed of free charged particles whose motion is influenced by the electromagnetic fields as well as by the externally applied magnetic field. As a consequence, the electromagnetic response of a magnetized plasma cannot be described by a scalar permittivity. Instead, a more general constitutive model is required to account for its anisotropic and dispersive behaviour [15].

In the frequency range considered in this work, the plasma response is described using the cold plasma approximation, which provides an accurate representation of electromagnetic wave propagation while considerably simplifying the governing equations.

The cold plasma approximation relies on the following assumptions [15]:

- the plasma is treated as a continuous fluid;
- the characteristic thermal velocity is assumed to be small compared with the phase velocity of the wave;
- the electromagnetic perturbations are assumed to be small, allowing a linear response around an equilibrium state;
- the equilibrium magnetic field is prescribed and time-independent.

Under these assumptions, the interaction between the electromagnetic field and the plasma can be incorporated into Maxwell's equations through an effective dielectric tensor. Consequently, the scalar relative permittivity introduced in section 2.3, is replaced by a frequency-dependent tensor describing the anisotropic electromagnetic response of the plasma.

The neglect of thermal effects can be described through the condition:

$$v_{th} \ll v_{ph}. \tag{2.32}$$

where v_{th} denotes the characteristic thermal velocity of the plasma particles, and v_{ph} the phase velocity of the electromagnetic wave. Under this assumption, thermal pressure effects can be neglected, and the fluid stress tensor vanishes,

$$\mathbf{\Phi}_j = \mathbf{0}, \tag{2.33}$$

where the subscript j denotes the plasma species (typically electrons or ions). The dynamics of each plasma species are then governed by the fluid momentum equation [15],

$$m_e \left(\frac{\partial \mathbf{v}_e}{\partial t} + \mathbf{v}_e \cdot \nabla \mathbf{v}_e \right) = -e \left(\mathbf{E} + \mathbf{v}_e \times \mathbf{B} \right) - \nabla \cdot \mathbf{\Phi}_e, \tag{2.34}$$

$$m_i \left(\frac{\partial \mathbf{v}_i}{\partial t} + \mathbf{v}_i \cdot \nabla \mathbf{v}_i \right) = e \left(\mathbf{E} + \mathbf{v}_i \times \mathbf{B} \right) - \nabla \cdot \mathbf{\Phi}_i. \tag{2.35}$$

Assuming small harmonic perturbations about an equilibrium state, the fluid equations can be linearized and solved analytically. The resulting expressions for the particle velocities describe the response of each plasma species to the applied electromagnetic field and constitute the basis for the derivation of the plasma current density. The resulting velocity components for a charged species are given by [15]

$$v_x = \frac{q}{m} \frac{i\omega E_x - \omega_c E_y}{\omega^2 - \omega_c^2}, \tag{2.36}$$

$$\tag{2.37}$$

$$v_y = \frac{q}{m} \frac{i\omega E_y + \omega_c E_x}{\omega^2 - \omega_c^2}, \tag{2.38}$$

$$\tag{2.39}$$

$$v_z = \frac{iq}{m\omega} E_z, \tag{2.40}$$

where q and m denote the charge and mass of the plasma species under consideration, respectively, and

$$\omega_c = \frac{qB_0}{m} \tag{2.41}$$

is the cyclotron frequency associated with the equilibrium magnetic field B_0.

The velocity expressions above describe the linear response of each plasma species to the applied electromagnetic field. They can therefore be used to compute the plasma current density, which is needed to provide the constitutive relation required to derive the cold-plasma dielectric tensor.

$$\mathbf{J} = en_0(\mathbf{v}_i - \mathbf{v}_e) = \boldsymbol{\sigma}\mathbf{E}, \tag{2.42}$$

where n_0 is the equilibrium plasma density. Since the particle velocities are linear functions of the electric field, the current density can be expressed in the constitutive form

$$\mathbf{J} = \boldsymbol{\sigma}\mathbf{E}, \tag{2.43}$$

where $\boldsymbol{\sigma}$ is the plasma conductivity tensor.

Substituting the constitutive relation into Ampère's law,

$$\nabla \times \mathbf{H} = \mathbf{J} - i\omega\mathbf{D}, \tag{2.44}$$

and recalling that

$$\mathbf{D} = \varepsilon_0\mathbf{E}, \tag{2.45}$$

the displacement field can be written as

$$\mathbf{D} = \varepsilon_0 \left(\mathbf{I} + \frac{i}{\omega\varepsilon_0}\boldsymbol{\sigma}\right)\mathbf{E}, \tag{2.46}$$

which defines the cold-plasma dielectric tensor through

$$\mathbf{D} = \varepsilon_0\boldsymbol{\varepsilon}\mathbf{E}. \tag{2.47}$$

The linear relation between the plasma current density and the electric field can be incorporated into Maxwell's equations through an effective dielectric tensor. Choosing a Cartesian coordinate system in which the equilibrium magnetic field is aligned with the z-axis,

$$\mathbf{B}_0 = B_0\hat{\mathbf{z}}, \tag{2.48}$$

the relative permittivity tensor of a collisionless cold magnetized plasma can be expressed in Stix notation as [15]

$$\boldsymbol{\varepsilon}_{\mathrm{cp}} = \begin{pmatrix} \varepsilon_\perp & -i\varepsilon_{xy} & 0 \\ i\varepsilon_{xy} & \varepsilon_\perp & 0 \\ 0 & 0 & \varepsilon_\parallel \end{pmatrix} = \begin{pmatrix} S & -iD & 0 \\ iD & S & 0 \\ 0 & 0 & P \end{pmatrix}. \tag{2.49}$$

The diagonal coefficients S and P describe the plasma response perpendicular and parallel to the equilibrium magnetic field, respectively, whereas the off-diagonal coefficient D accounts for the gyrotropic response introduced by the magnetic field. These coefficients are defined in terms of the standard Stix parameters [15] as

$$S = \frac{R+L}{2} = 1 - \sum_j \frac{\omega_{pj}^2}{\omega^2 - \omega_{cj}^2}, \tag{2.50a}$$

$$D = \frac{R-L}{2} = \sum_j \frac{\omega_{cj}\omega_{pj}^2}{\omega\left(\omega^2 - \omega_{cj}^2\right)}, \tag{2.50b}$$

$$P = 1 - \sum_j \frac{\omega_{pj}^2}{\omega^2}, \tag{2.50c}$$

$$R = S + D = 1 - \sum_j \frac{\omega_{pj}^2}{\omega\left(\omega + \omega_{cj}\right)}, \tag{2.50d}$$

$$L = S - D = 1 - \sum_j \frac{\omega_{pj}^2}{\omega\left(\omega - \omega_{cj}\right)}. \tag{2.50e}$$

Here, the summation is performed over all charged plasma species. The plasma frequency of species j is defined as

$$\omega_{pj}^2 = \frac{n_j q_j^2}{\varepsilon_0 m_j}, \tag{2.51}$$

where n_j, q_j, and m_j denote its particle density, charge, and mass, respectively. The quantity ω_{cj} is the signed cyclotron frequency introduced in eq. (2.41), here written explicitly for species j as

$$\omega_{cj} = \frac{q_j B_0}{m_j}. \tag{2.52}$$

Equivalently, the particle charge may be written as $q_j = Z_j e$, where Z_j is the signed charge number and e is the elementary charge.

The coefficients R and L represent the dielectric response to right-hand and left-hand circularly polarized waves, respectively [15]. Since the cyclotron frequency is signed, electrons and positively charged ions contribute differently to the off-diagonal component D. This coupling between the two perpendicular electric-field components is responsible for the anisotropic character of wave propagation in a magnetized plasma.

Substituting eq. (2.49) for the scalar relative permittivity in eq. (2.20) gives the vector wave equation used to describe electromagnetic propagation in the cold-plasma model.

The Stix tensor completely characterizes the dielectric response of a collisionless cold magnetized plasma. Unlike conventional isotropic materials, where the relative permittivity is represented by a scalar quantity, the electromagnetic properties of a magnetized plasma depend on both the wave frequency and the direction of propagation with respect to the equilibrium magnetic field.

Consequently, the scalar relative permittivity ε_r appearing in the vector wave equation derived in the previous section is replaced by the cold-plasma dielectric tensor,

$$\varepsilon_r \longrightarrow \boldsymbol{\varepsilon}_{\mathrm{cp}}, \tag{2.53}$$

thus, providing a mathematical description of electromagnetic wave propagation in magnetized fusion plasmas. The governing equation for the electric field therefore becomes

$$\nabla \times \left(\mu_r^{-1} \nabla \times \mathbf{E}\right) - k_0^2 \boldsymbol{\varepsilon}_{\mathrm{cp}} \mathbf{E} = -ik_0 \eta_0 \mathbf{J}. \tag{2.54}$$

In many practical applications, however, electromagnetic waves also interact with conductive materials, whose effect can be incorporated through a modified constitutive relation, as discussed in section 2.5.

Although the cold-plasma approximation provides an accurate description of the propagation properties discussed above, it does not account for finite-temperature effects and wave–particle interactions [15, 16]. Their description requires a hot-plasma dielectric tensor, which lies beyond the scope of the present work.

2.5 Conductive Media

In practical fusion devices, electromagnetic waves propagate not only through the plasma but also through surrounding conductive components, such as antennas, vacuum vessel structures, and plasma-facing components. The finite electrical conductivity of these materials modifies their electromagnetic response and must therefore be incorporated into the constitutive relations.

For an isotropic conductive medium, the conduction current density is described by Ohm's law [14],

$$\mathbf{J}_c = \sigma \mathbf{E}, \tag{2.55}$$

where σ denotes the electrical conductivity of the material and is expressed in S/m.

The total current density appearing in Ampère's law can therefore be written as

$$\mathbf{J} = \mathbf{J}_s + \mathbf{J}_c = \mathbf{J}_s + \sigma \mathbf{E}, \tag{2.56}$$

where $\mathbf{J}_s$ represents the impressed source current density.

Substituting this expression into the frequency-domain form of Ampère's law,

$$\nabla \times \mathbf{H} = \mathbf{J}_s + \sigma \mathbf{E} - i\omega \mathbf{D}, \tag{2.57}$$

and using the constitutive relation

$$\mathbf{D} = \varepsilon_0 \boldsymbol{\varepsilon}_r \mathbf{E}, \tag{2.58}$$

yields

$$\nabla \times \mathbf{H} = \mathbf{J}_s - i\omega\varepsilon_0 \left(\boldsymbol{\varepsilon}_r + i\frac{\sigma}{\omega\varepsilon_0}\mathbf{I} \right) \mathbf{E}. \tag{2.59}$$

For each material, the conductive contribution is incorporated into its corresponding relative permittivity tensor according to [14],

$$\boldsymbol{\varepsilon}_{r,\text{eff}}^{(m)} = \boldsymbol{\varepsilon}_r^{(m)} + i\frac{\sigma^{(m)}}{\omega\varepsilon_0}\mathbf{I}, \tag{2.60}$$

where m denotes the material considered. For an isotropic dielectric, $\boldsymbol{\varepsilon}_r^{(m)} = \varepsilon_r^{(m)}\mathbf{I}$, whereas for a cold-plasma material it corresponds to the cold-plasma dielectric tensor introduced in section 2.4. Since the conductivity considered here is isotropic, its contribution is proportional to $\mathbf{I}$ and therefore modifies only the diagonal components of the permittivity tensor of that material.

In the vector wave equation 2.20, the relative permittivity associated with each material is therefore replaced by its corresponding effective value $\boldsymbol{\varepsilon}_{r,\text{eff}}^{(m)}$.

2.6 Wave Propagation in Magnetized Plasmas

Having introduced the cold-plasma dielectric tensor, its consequences for electromagnetic wave propagation can now be examined. Depending on the wave frequency, polarization, and propagation direction with respect to the equilibrium magnetic field, different wave modes can arise, characterized by different dispersion and resonance properties [15, 16]. These properties are particularly relevant to radio-frequency plasma heating, including Ion Cyclotron Resonance Heating (ICRH) [7].

To analyse the propagation characteristics, it is convenient to introduce the refractive index vector [15],

$$\mathbf{n} = \frac{c}{\omega}\mathbf{k}, \tag{2.61}$$

where $\mathbf{k}$ is the wavevector and c is the speed of light in vacuum.

The refractive index can be decomposed into components parallel and perpendicular to the equilibrium magnetic field,

$$\mathbf{n} = \mathbf{n}_\perp + \mathbf{n}_\parallel. \tag{2.62}$$

Throughout this work, wave propagation is assumed to take place in the plane perpendicular to the equilibrium magnetic field, corresponding to the two-dimensional geometry solved by the `EMWAVE` module. Consequently,

$$n_\parallel = 0, \tag{2.63}$$

or equivalently,

$$\theta = \frac{\pi}{2}, \tag{2.64}$$

where θ denotes the angle between the wavevector $\mathbf{k}$ and the equilibrium magnetic field $\mathbf{B}_0$.

For the two-dimensional configuration considered in this work, the propagation properties of the TE and TM polarizations introduced in section 2.3, can be analyzed by considering plane-wave solutions of the Maxwell's equations in the cold plasma medium. For perpendicular propagation, this leads to two indipendent modes, conventionally refferred to as Ordinary (O) and Extraordinary (X) waves [15].

Substituting the cold-plasma dielectric tensor 2.49 into Maxwell's equations (2.1)–(2.4), and considering plane-wave solutions leads to the following homogeneous system for the electric-field components [15]:

$$\begin{pmatrix} S - n^2\cos^2\theta & -iD & n^2\cos\theta\sin\theta \\ iD & S - n^2 & 0 \\ n^2\cos\theta\sin\theta & 0 & P - n^2\sin^2\theta \end{pmatrix} \begin{pmatrix} E_x \\ E_y \\ E_z \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}. \tag{2.65}$$

Non-trivial solutions of eq. (2.65) exist only when the determinant of the coefficient matrix vanishes. For the perpendicular propagation considered in this work, the system reduces to

$$\begin{pmatrix} S & -iD & 0 \\ iD & S-n^2 & 0 \\ 0 & 0 & P-n^2 \end{pmatrix} \begin{pmatrix} E_x \\ E_y \\ E_z \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}. \tag{2.66}$$

2.6.1 Ordinary Wave

For perpendicular propagation, the third row of eq. (2.66) is decoupled from the remaining system and gives

$$\left(P - n^2\right) E_z = 0. \tag{2.67}$$

For a non-trivial solution with $E_z \neq 0$, the dispersion relation of the ordinary mode is therefore [15]

$$n_O^2 = P = 1 - \sum_j \frac{\omega_{pj}^2}{\omega^2}. \tag{2.68}$$

The electric field of the ordinary wave is parallel to the equilibrium magnetic field, while the wavevector is perpendicular to it [15]. In the coordinate system adopted here, this corresponds to

$$\mathbf{E} = E_z \hat{\mathbf{z}}, \qquad \mathbf{k} \perp \mathbf{B}_0. \tag{2.69}$$

Since the coefficient P does not contain the cyclotron frequency, the dispersion relation of the ordinary wave is independent of the equilibrium magnetic-field magnitude. The plasma magnetization therefore does not directly affect this mode.

Propagation requires a real refractive index, and hence

$$n_O^2 > 0. \tag{2.70}$$

Using eq. (2.68), the ordinary wave encounters a cutoff when

$$P = 0, \tag{2.71}$$

which, for a single-species plasma, corresponds to

$$\omega = \omega_p, \tag{2.72}$$

where ω_p is the plasma frequency, introduced in eq. (2.51).

Because the mass of the electrons is much lower than that of the ions ($m_e \ll m_i$), the electron plasma frequency is generally much larger than the ion plasma frequency and therefore dominates the cutoff condition. Consequently, the ordinary-wave cutoff is commonly approximated by

$$\omega \simeq \omega_{pe}, \tag{2.73}$$

$$\omega_{pe} = \sqrt{\frac{n_e e^2}{m_e \varepsilon_0}}. \tag{2.74}$$

As the wave propagates from the plasma edge towards the core, the electron density progressively increases. Because of this, the parameter P decreases as the wave penetrates into the plasma. When the condition

$$\omega = \omega_{pe} \tag{2.75}$$

is satisfied, the ordinary wave reaches its cutoff layer, where

$$n_O^2 = 0. \tag{2.76}$$

Beyond this point, n_O^2 becomes negative and the refractive index is purely imaginary. Consequently, the wave becomes evanescent and its amplitude decays exponentially, preventing further propagation towards the plasma core [15].

For this reason, the ordinary wave is generally unable to deliver electromagnetic power to the plasma core. The extraordinary wave, discussed in the following section 2.6.2, exhibits substantially different propagation characteristics that make it suitable for Ion Cyclotron Resonance Heating.

2.6.2 Extraordinary Wave

The extraordinary (X) wave is obtained from the coupled system involving the transverse electric-field components. For a non-trivial solution, the determinant of the corresponding 2×2 system must vanish, leading to the dispersion relation [15]

$$n_X^2 = \frac{RL}{S}, \tag{2.77}$$

where the Stix parameters R, L, and S were introduced in section 2.4.

For a plasma composed of electrons and a single ion species, the dispersion relation can be written explicitly as

$$n_X^2 = \frac{RL}{S} = \frac{\left[(\omega+\omega_{ci})(\omega-\omega_{ce})-\omega_p^2\right]\left[(\omega-\omega_{ci})(\omega+\omega_{ce})-\omega_p^2\right]}{(\omega^2-\omega_{ci}^2)(\omega^2-\omega_{ce}^2)+\omega_p^2(\omega_{ce}\omega_{ci}-\omega^2)}, \tag{2.78}$$

where

$$\omega_p^2 = \omega_{pe}^2 + \omega_{pi}^2,$$

and ω_{cj} denotes the cyclotron frequency of species j.

Unlike the ordinary wave, the extraordinary mode is directly affected by the equilibrium magnetic field through the cyclotron frequencies. Consequently, its propagation depends simultaneously on the plasma density, the magnetic field strength, and the operating frequency.

The resonance conditions correspond to the singularities of eq. (2.78), that is, when the denominator vanishes. These occur at

$$\omega^2 = \frac{\omega_e^2+\omega_i^2}{2} \pm \sqrt{\left(\frac{\omega_e^2-\omega_i^2}{2}\right)^2 + \omega_{pe}^2\omega_{pi}^2}, \tag{2.79}$$

where

$$\omega_j^2 = \omega_{pj}^2 + \omega_{cj}^2.$$

Since the electron mass is much smaller than the ion mass ($m_e \ll m_i$), these expressions simplify to the well-known upper and lower hybrid resonance frequencies [15, 16],

$$\omega_{UH}^2 \simeq \omega_{pe}^2 + \omega_{ce}^2, \tag{2.80}$$

$$\omega_{LH}^2 \simeq \omega_{ce}\omega_{ci}\left(\frac{\omega_{pe}^2+\omega_{ce}\omega_{ci}}{\omega_{pe}^2+\omega_{ce}^2}\right). \tag{2.81}$$

Similarly, the cutoff frequencies are obtained from the zeros of the numerator of eq. (2.78),

$$\omega_X = \sqrt{\left(\frac{\omega_{ce}+\omega_{ci}}{2}\right)^2 + \omega_p^2} \pm \frac{\omega_{ce}-\omega_{ci}}{2}. \tag{2.82}$$

These cutoff frequencies define the limits beyond which the wave can no longer propagate inside the plasma [15]. They are commonly referred to as the *left-hand cutoff* ($L=0$) and the *right-hand cutoff* ($R=0$).

The presence of cutoff and resonance layers plays a fundamental role in the accessibility of radio-frequency waves. Depending on the plasma parameters and the injected wave frequency, an electromagnetic wave may be reflected, become evanescent, tunnel through a cutoff region, or undergo mode conversion [15, 16]. Accurately capturing these phenomena is essential for the numerical modelling of Ion Cyclotron Resonance Heating (ICRH).

Unlike the ordinary wave, the extraordinary mode is able to propagate through regions of the plasma that are inaccessible to the O-mode, because of its different dispersion characteristics and its interaction with the equilibrium magnetic field. As a result, the X-wave can reach the ion cyclotron resonance layer located deeper inside the plasma, where efficient transfer of electromagnetic energy to the ions takes place.

For this reason, the extraordinary wave constitutes the primary electromagnetic mode employed in Ion Cyclotron Resonance Heating (ICRH) systems. Accurately describing its propagation, including the presence of cutoff and resonance layers, is therefore essential for predicting the deposition of radio-frequency power in magnetically confined fusion plasmas.

3 Finite Element Method

The governing equations presented in Chapter 2 generally do not admit analytical solutions for the complex geometries encountered in fusion devices. Numerical techniques are therefore required to approximate the electromagnetic fields.

Among the various numerical methods available, the Finite Element Method (FEM) is particularly well suited for the solution of electromagnetic problems in complex geometries [8]. It is a numerical tool capable of solving linear and non-linear partial differential equations (PDEs), with a set of boundary conditions. Such type of problem is generally referred to as boundary value problem (BVP). The ability of FEM to employ unstructured meshes, accurately represent curved boundaries, and naturally handle heterogeneous material properties makes it especially useful for applications involving fusion plasmas and radio-frequency antennas.

The `EMWAVE` module of the Alya multiphysics code employs the finite element method to discretize the vector wave equation derived in the previous sections. The main concepts of the FEM formulation adopted in this work are briefly reviewed below.

3.1 Theoretical Foundations of FEM

The fundamental idea of the finite element method consists in subdividing the computational domain into a finite number of non-overlapping subdomains, called *finite elements* [8]. The collection of these elements forms the computational mesh, over which the governing differential equations are solved. Each element is defined by a set of nodes whose number depends on both the element topology and the interpolation order. Typical element shapes include triangles and quadrilaterals in two dimensions, and tetrahedra and hexahedra in three dimensions, as illustrated in fig. 2a. Furthermore, an element may employ either first-order (linear), second-order (quadratic), or higher-order interpolation functions, as shown in fig. 2b.

Within every element, the unknown field is approximated by a linear combination of interpolation functions, usually referred to as *shape functions* [8],

$$\mathbf{u}(\mathbf{x}) = \sum_{i=1}^{N} u_i N_i(\mathbf{x}), \tag{3.1}$$

where N_i denotes the shape function associated with the i-th node, u_i is the corresponding degree of freedom, and N is the number of nodes belonging to the element.

The shape functions interpolate the unknown field inside the element from its nodal values. Consequently, the accuracy of the numerical approximation depends strongly on both the interpolation order and the quality of the mesh.

3.2 Galerkin Finite Element Formulation

The finite element method seeks an approximate solution to a boundary-value problem of the general form

$$Lu = f, \tag{3.2}$$

where L is a differential operator, f is a known source term, and u denotes the unknown field.

Examples of boundary-value problems encountered in mathematical physics include the transport equation, the heat equation, the wave equation, and Laplace's equation.

Since the computational domain is decomposed into finite elements, eq. (3.2) is solved locally over each element,

$$Lu^e = f^e. \tag{3.3}$$

The exact solution is generally unknown and is therefore approximated using the shape functions introduced previously,

$$u^e = \sum_{j=1}^{n} u_j^e N_j^e, \tag{3.4}$$

where n denotes the number of nodes of the element.

Substituting the approximate solution given by eq. (3.4) into the governing equation does not satisfy eq. (3.2) exactly. The resulting error is called the *residual* [8],

$$R^e = Lu^e - f^e. \tag{3.5}$$

Rather than enforcing the residual to vanish pointwise, the finite element method requires it to vanish in an average sense over the element domain. This idea constitutes the basis of the *weighted residual method* [8], which imposes

$$\int_{\Omega^e} R^e \, w_i^e \, d\Omega = 0, \tag{3.6}$$

where w_i^e denotes a suitable weighting function.

Different choices of the weighting functions lead to different weighted residual methods. The most common are

- Collocation method;
- Subdomain collocation method;
- Galerkin method;
- Petrov–Galerkin method;
- Least-squares method.

In the present work, the Galerkin method is adopted. In this approach, the weighting functions are chosen to coincide with the shape functions [8],

$$w_i^e = N_i^e. \tag{3.7}$$

The Galerkin formulation is widely used in finite-element methods because it provides a systematic framework for deriving the weak form while using the same approximation space for the trial and weighting functions [8].

Substituting the finite element approximation into eq. (3.6) yields

$$\int_{\Omega^e} w_i^e L \left(\sum_{j=1}^{n} u_j^e N_j^e \right) d\Omega = \int_{\Omega^e} w_i^e f^e \, d\Omega. \tag{3.8}$$

Applying the Galerkin choice,

$$\sum_{j=1}^{n} u_j^e \left(\int_{\Omega^e} N_i^e L N_j^e \, d\Omega \right) = \int_{\Omega^e} N_i^e f^e \, d\Omega. \tag{3.9}$$

from which one can identify a linear system of equations $\mathbf{Au} = \mathbf{b}$, where

$$A_j^i = \int_{\Omega^e} N_i^e L N_j^e \mathrm{d}\Omega \tag{3.10}$$

$$b^i = \int_{\Omega^e} N_i^e f^e \mathrm{d}\Omega \tag{3.11}$$

3.3 Element Categories

Finite elements can be classified according to their spatial dimension, geometrical shape, and interpolation order [8]. The developments concerning new element types presented in this work focus on two-dimensional triangular and quadrilateral discretizations.

The nodes of an element have arbitrary coordinates in the physical domain, denoted by (x, y). However, finite-element interpolation is more conveniently defined over a standard reference geometry, commonly referred to as the *master element* or *reference element*, described by the local coordinates (ξ, η) [8].

A mapping relates each physical element to the corresponding master element. This approach allows the shape functions and numerical integration rules to be defined once on a standard domain and subsequently reused for every element of the same type in the mesh.

Within each element, the approximate solution is expressed as

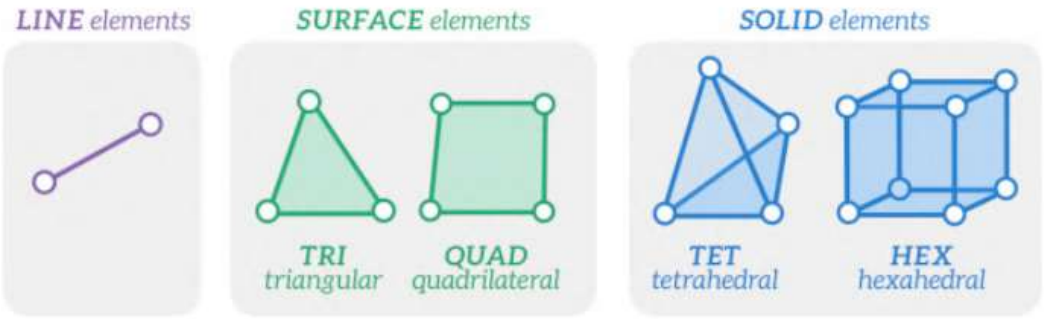

(a) Examples of finite elements in one, two, and three dimensions.

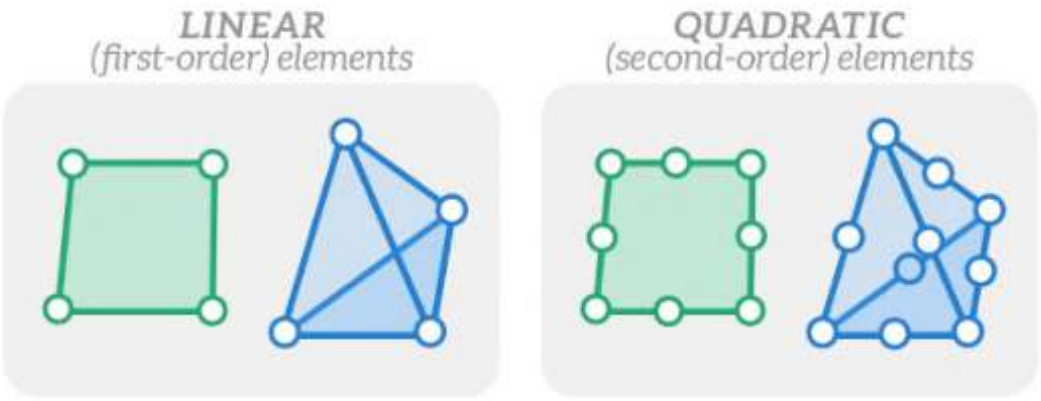

(b) Comparison between linear and quadratic finite elements.

Figure 2: Schematic representation of finite-element discretizations and interpolation orders. Extracted from [10].

$$u^e(\xi,\eta) = \sum_{j=1}^{n} u_j^e N_j(\xi,\eta), \tag{3.12}$$

where u_j^e denotes the value of the unknown associated with node j, N_j is the corresponding shape function, and n is the number of nodes of the element.

The shape functions satisfy the Kronecker-delta interpolation property,

$$N_j(\xi_i,\eta_i) = \delta_{ij} = \begin{cases} 1, & i = j, \\ 0, & i \neq j, \end{cases} \tag{3.13}$$

which ensures that the finite-element approximation reproduces the nodal values exactly. In addition, the shape functions form a partition of unity,

$$\sum_{j=1}^{n} N_j(\xi,\eta) = 1. \tag{3.14}$$

3.3.1 Linear Triangular Element

A linear triangular element contains three nodes, one at each vertex, and possesses straight edges. The nodes are conventionally numbered in the counter-clockwise direction [8].

For the reference triangle, the nodal coordinates are chosen as

$$(\xi_1,\eta_1) = (0;0), \qquad (\xi_2,\eta_2) = (1;0), \qquad (\xi_3,\eta_3) = (0;1). \tag{3.15}$$

Since the interpolation is linear, the approximate field can be written as

$$u^e(\xi,\eta) = c_1 + c_2\xi + c_3\eta. \tag{3.16}$$

Imposing the nodal interpolation conditions of eq. (3.15) to eq. (3.16), leads to

$$c_1 = u_1^e \tag{3.17}$$
$$c_2 = -u_1^e + u_2^e \tag{3.18}$$
$$c_3 = -u_1^e + u_3^3. \tag{3.19}$$

Substituting these coefficients in equation eq. (3.16), yields

$$u^e(\xi,\eta) = u_1^e(1-\xi-\eta) + u_2^e\xi + u_3^e\eta. \tag{3.20}$$

Comparing with eq. (3.12), it is easily found that:

$$N_1(\xi,\eta) = 1-\xi-\eta, \tag{3.21a}$$
$$N_2(\xi,\eta) = \xi, \tag{3.21b}$$
$$N_3(\xi,\eta) = \eta. \tag{3.21c}$$

The physical coordinates are interpolated using the same shape functions [8],

$$x(\xi,\eta) = \sum_{j=1}^{3} N_j(\xi,\eta)x_j, \tag{3.22a}$$
$$y(\xi,\eta) = \sum_{j=1}^{3} N_j(\xi,\eta)y_j. \tag{3.22b}$$

The transformation between the master and physical domains is described by the Jacobian matrix,

$$\mathbf{J} = \begin{pmatrix} \dfrac{\partial x}{\partial \xi} & \dfrac{\partial x}{\partial \eta} \\ \dfrac{\partial y}{\partial \xi} & \dfrac{\partial y}{\partial \eta} \end{pmatrix}. \tag{3.23}$$

The derivatives of the shape functions with respect to the physical coordinates are obtained through

$$\begin{pmatrix} \dfrac{\partial N_j}{\partial x} \\ \dfrac{\partial N_j}{\partial y} \end{pmatrix} = \mathbf{J}^{-T} \begin{pmatrix} \dfrac{\partial N_j}{\partial \xi} \\ \dfrac{\partial N_j}{\partial \eta} \end{pmatrix}. \tag{3.24}$$

3.3.2 Quadratic Triangular Element

The accuracy of the finite element approximation can be improved by increasing the interpolation order. In the case of triangular elements, this is achieved by introducing one additional node at the midpoint of each edge, leading to a six-node quadratic element [8], as shown in fig. 2b.

The reference element remains the same as for the linear triangle, but three additional nodes are located at the edge midpoints,

$$(\xi_4, \eta_4) = \left(\frac{1}{2}, 0\right), \qquad (\xi_5, \eta_5) = \left(\frac{1}{2}, \frac{1}{2}\right), \qquad (\xi_6, \eta_6) = \left(0, \frac{1}{2}\right). \tag{3.25}$$

The field inside the element is now approximated by a second-order polynomial,

$$u^e(\xi, \eta) = c_1 + c_2\xi + c_3\eta + c_4\xi\eta + c_5\xi^2 + c_6\eta^2. \tag{3.26}$$

Imposing the interpolation conditions at the six nodes of the reference element leads to the quadratic Lagrange shape functions [8],

$$N_1 = (1 - \xi - \eta)(1 - 2\xi - 2\eta), \tag{3.27a}$$
$$N_2 = \xi(2\xi - 1), \tag{3.27b}$$
$$N_3 = \eta(2\eta - 1), \tag{3.27c}$$
$$N_4 = 4\xi(1 - \xi - \eta), \tag{3.27d}$$
$$N_5 = 4\xi\eta, \tag{3.27e}$$
$$N_6 = 4\eta(1 - \xi - \eta). \tag{3.27f}$$

The finite-element approximation therefore becomes

$$u^e(\xi, \eta) = \sum_{j=1}^{6} u_j^e N_j(\xi, \eta). \tag{3.28}$$

As for the linear element, the physical coordinates are interpolated using the same shape functions,

$$x(\xi, \eta) = \sum_{j=1}^{6} N_j(\xi, \eta) x_j, \tag{3.29a}$$
$$y(\xi, \eta) = \sum_{j=1}^{6} N_j(\xi, \eta) y_j. \tag{3.29b}$$

The use of quadratic interpolation significantly improves the representation of both the computational geometry and the solution field. In particular, curved boundaries can be approximated much more accurately than with first-order elements, while the additional interpolation freedom generally reduces the discretization error for smooth solutions [8].

The improved accuracy, however, comes at the expense of an increased number of degrees of freedom per element, larger element matrices, and a higher computational cost during both assembly and solution of the resulting linear system [8]. Consequently, the choice between linear and quadratic elements represents a compromise between accuracy and computational efficiency. The implementation of quadratic triangular elements within the nodal formulation of the `EMWAVE` module constitutes one of the developments carried out during this internship and is discussed in detail in Chapter 4.

3.3.3 Bilinear Quadrilateral Element

An alternative to triangular discretizations is provided by quadrilateral elements. Owing to their tensor-product structure, quadrilateral elements often achieve a higher accuracy per degree of freedom in regions with regular geometries and aligned meshes [8].

The reference quadrilateral element is defined over the square

$$-1 \leq \xi \leq 1, \qquad -1 \leq \eta \leq 1, \tag{3.30}$$

with the four corner nodes located at

$$(\xi_1, \eta_1) = (-1, -1), \qquad (\xi_2, \eta_2) = (1, -1), \tag{3.31}$$

$$(\xi_3, \eta_3) = (1, 1), \qquad (\xi_4, \eta_4) = (-1, 1). \tag{3.32}$$

Unlike triangular elements, whose interpolation is expressed in terms of complete polynomials, quadrilateral elements are constructed as tensor products of one-dimensional interpolation functions. Consequently, the field inside the element is approximated by the bilinear polynomial [8]

$$u^e(\xi, \eta) = c_1 + c_2\xi + c_3\eta + c_4\xi\eta. \tag{3.33}$$

Applying the interpolation conditions at the four nodes of the master element yields the bilinear shape functions [8],

$$N_1(\xi, \eta) = \frac{1}{4}(1 - \xi)(1 - \eta), \tag{3.34a}$$

$$N_2(\xi, \eta) = \frac{1}{4}(1 + \xi)(1 - \eta), \tag{3.34b}$$

$$N_3(\xi, \eta) = \frac{1}{4}(1 + \xi)(1 + \eta), \tag{3.34c}$$

$$N_4(\xi, \eta) = \frac{1}{4}(1 - \xi)(1 + \eta). \tag{3.34d}$$

The finite-element approximation is therefore

$$u^e(\xi, \eta) = \sum_{j=1}^{4} u_j^e N_j(\xi, \eta). \tag{3.35}$$

As in the case of triangular elements, the physical coordinates are interpolated using the same shape functions,

$$x(\xi, \eta) = \sum_{j=1}^{4} N_j(\xi, \eta) x_j, \tag{3.36a}$$

$$y(\xi, \eta) = \sum_{j=1}^{4} N_j(\xi, \eta) y_j. \tag{3.36b}$$

The mapping between the master and physical domains is again described by the Jacobian matrix introduced previously. Since the mapping is generally nonlinear, distorted quadrilateral elements can accurately represent complex geometries while preserving the same interpolation framework.

The original nodal formulation of the `EMWAVE` module only supported triangular finite elements. One of the main contributions of the present work was the implementation of bilinear quadrilateral elements within the nodal solver, extending the range of meshes that can be employed for electromagnetic simulations. The numerical implementation and its validation are presented in Chapter 4.

3.4 Weak Formulation

Starting from the general scalar wave equation introduced in eq. (2.25), the finite-element formulation is constructed by considering its restriction to an element Ω^e. Accordingly,

$$-\nabla \cdot (\mathbf{p}^e \cdot \nabla u^e) + q^e u^e = f^e, \tag{3.37}$$

where the superscript e denotes quantities restricted to the element.

As discussed in section 3.2, the finite element method does not require the differential equation to be satisfied exactly at every point of the element. Instead, the weighted residual method imposes that the residual vanishes in an integral sense [8]. Therefore,

$$\iiint_{\Omega^e} w_i^e \left[-\nabla \cdot (\mathbf{p}^e \cdot \nabla u^e) + q^e u^e - f^e\right] dV = 0. \tag{3.38}$$

The first term in eq. (3.38) contains second-order spatial derivatives of the unknown field. In order to reduce the order of these derivatives, the vector identity

$$\nabla \cdot (\psi \mathbf{A}) = \psi \nabla \cdot \mathbf{A} + \mathbf{A} \cdot \nabla \psi \tag{3.39}$$

is employed, where ψ is a scalar field and $\mathbf{A}$ is a vector field.

Taking

$$\psi = w_i^e, \qquad \mathbf{A} = \mathbf{p}^e \cdot \nabla u^e, \tag{3.40}$$

one obtains

$$w_i^e \nabla \cdot (\mathbf{p}^e \cdot \nabla u^e) = \nabla \cdot [w_i^e \, (\mathbf{p}^e \cdot \nabla u^e)] - (\mathbf{p}^e \cdot \nabla u^e) \cdot \nabla w_i^e. \tag{3.41}$$

Substituting eq. (3.41) into eq. (3.38) gives

$$\begin{aligned} \iiint_{\Omega^e} \Big[& (\mathbf{p}^e \cdot \nabla u^e) \cdot \nabla w_i^e - \nabla \cdot (w_i^e \mathbf{p}^e \cdot \nabla u^e) \\ & + q^e w_i^e u^e - w_i^e f^e \Big] \, dV = 0. \end{aligned} \tag{3.42}$$

The divergence theorem can now be applied to the second term,

$$\iiint_{\Omega^e} \nabla \cdot \mathbf{F} \, dV = \iint_{\Gamma^e} \mathbf{F} \cdot \mathbf{n} \, dS, \tag{3.43}$$

where Γ^e denotes the boundary of the element and $\mathbf{n}$ is the outward unit normal vector. Consequently, eq. (3.42) becomes

$$\begin{aligned} & \iiint_{\Omega^e} \left[(\mathbf{p}^e \cdot \nabla u^e) \cdot \nabla w_i^e + q^e w_i^e u^e \right] dV \\ & - \iint_{\Gamma^e} w_i^e \, (\mathbf{p}^e \cdot \nabla u^e) \cdot \mathbf{n} \, dS = \iiint_{\Omega^e} w_i^e f^e \, dV. \end{aligned} \tag{3.44}$$

Equation eq. (3.44) represents the weak formulation over a three-dimensional element. Compared with the strong form, the highest-order spatial derivative has been reduced from second to first order. This constitutes one of the main advantages of the finite element formulation, since less restrictive continuity requirements are imposed on the interpolation functions [8].

For the two-dimensional formulation considered in the nodal implementation of `EMWAVE`, invariance is assumed along the out-of-plane direction. The volume integrals can therefore be reduced to surface integrals over the two-dimensional element domain, while the surface integral over the three-dimensional boundary becomes a line integral along the element boundary. The weak form then reads

$$\begin{aligned} & \iint_{\Omega^e} \left[(\mathbf{p}^e \cdot \nabla u^e) \cdot \nabla w_i^e + q^e w_i^e u^e \right] d\Omega \\ & - \int_{\Gamma^e} w_i^e \, (\mathbf{p}^e \cdot \nabla u^e) \cdot \mathbf{n} \, d\Gamma = \iint_{\Omega^e} w_i^e f^e \, d\Omega, \end{aligned} \tag{3.45}$$

where Ω^e denotes the two-dimensional element domain and $\Gamma^e = \partial\Omega^e$ its one-dimensional boundary.

For interior element boundaries, the corresponding boundary contributions are cancelled during the global assembly by the contribution of the neighbouring element, whose outward normal points in the opposite direction. The boundary integral therefore only needs to be explicitly retained on the external boundary of the computational domain, where the appropriate boundary conditions are imposed.

Adopting the Galerkin method [8],

$$w_i^e = N_i^e, \tag{3.46}$$

and introducing the finite element approximation (here reported for simplicity)

$$u^e = \sum_{j=1}^{n} u_j^e N_j^e, \tag{3.47}$$

eq. (3.45), neglecting the boundary contribution for an interior element, becomes

$$\sum_{j=1}^{n} u_j^e \left[\iint_{\Omega^e} \left(\left(\mathbf{p}^e \cdot \nabla N_j^e \right) \cdot \nabla N_i^e + q^e N_i^e N_j^e \right) d\Omega \right] = \iint_{\Omega^e} N_i^e f^e \, d\Omega. \tag{3.48}$$

The terms in eq. (3.48) can be directly identified with the coefficients of the elemental linear system,

$$\mathbf{A}^e \mathbf{u}^e = \mathbf{b}^e, \tag{3.49}$$

where the element matrix entries are

$$(A^e)_{ij} = \iint_{\Omega^e} \left[\left(\mathbf{p}^e \cdot \nabla N_j^e \right) \cdot \nabla N_i^e + q^e N_i^e N_j^e \right] d\Omega, \tag{3.50}$$

and the components of the elemental right-hand-side vector are

$$b_i^e = \iint_{\Omega^e} N_i^e f^e \, d\Omega. \tag{3.51}$$

Finally, the elemental systems are assembled according to the connectivity between the finite elements, giving rise to the global linear system [8]

$$\mathbf{A}\mathbf{u} = \mathbf{b}, \tag{3.52}$$

whose solution provides the finite element approximation of the electromagnetic field over the complete computational domain.

3.5 Edge Elements

The nodal finite element formulation presented in the previous sections has been successfully employed for many scalar partial differential equations, including the TE and TM formulations derived from Maxwell's equations. However, when solving the full vector wave equation, conventional nodal elements present several important limitations [8]. In a nodal discretization of a vector field, each field component is interpolated separately using scalar shape functions, resulting in n_{dim} scalar degrees of freedom per node. In contrast, edge elements employ vector-valued basis functions and associate one scalar degree of freedom with each mesh edge.

The electromagnetic field is governed by curl operators, and the physically relevant continuity condition across interfaces concerns the tangential component of the electric field rather than the field itself. Standard Lagrangian shape functions enforce the continuity of every field component at the mesh nodes, a constraint that is stronger than required by Maxwell's equations [8].

As a consequence, nodal formulations may generate non-physical solutions, commonly referred to as *spurious modes*, and may provide an inadequate representation of the curl operator [8]. These shortcomings become particularly important when solving three-dimensional vector electromagnetic problems.

To overcome these limitations, edge finite elements, originally introduced by Nédélec [12], employ vector-valued basis functions belonging to the $H(\text{curl})$ functional space. This construction naturally enforces the required continuity of the tangential field component across element interfaces [8].

The mathematical construction of the Nédélec basis functions, and the interpolation of the electric field are introduced in section 3.5.1.

The `EMWAVE` module supports both nodal and edge finite-element formulations. The former is used for the scalar TE and TM formulations described previously, whereas the latter allows the full vector wave equation to be discretized while preserving the tangential continuity required by the electromagnetic problem.

3.5.1 Nédélec Basis Functions

As discussed in section 3.5, the degrees of freedom of edge elements are associated with the mesh edges rather than with the mesh nodes. Consequently, the interpolation of the unknown field cannot be performed using scalar shape functions. Instead, vector-valued basis functions, known as *Nédélec basis functions*, are employed [12].

For triangular elements, the construction of the Nédélec basis functions can be conveniently expressed in terms of the barycentric coordinates [12]. The barycentric coordinates L_1, L_2, and L_3 are scalar functions associated with the three vertices of the reference triangle and satisfy

$$L_1 + L_2 + L_3 = 1. \tag{3.53}$$

Each barycentric coordinate satisfies the interpolation property

$$L_i = \begin{cases} 1, & \text{at vertex } i, \\ 0, & \text{at the other two vertices.} \end{cases} \tag{3.54}$$

For the reference triangular element introduced in section 3.3.1, the barycentric coordinates coincide with its first-order scalar shape functions and are therefore given by

$$L_1 = 1 - \xi - \eta, \qquad L_2 = \xi, \qquad L_3 = \eta. \tag{3.55}$$

For a first-order triangular edge element, one vector function is associated with each of its three edges. For an oriented edge connecting vertices i and j, the corresponding function is defined as [12]

$$\mathbf{W}_{ij} = L_i \nabla L_j - L_j \nabla L_i. \tag{3.56}$$

The orientation of the edges must be taken into account when constructing these functions [8]. Reversing the orientation changes their sign,

$$\mathbf{W}_{ji} = -\mathbf{W}_{ij}. \tag{3.57}$$

A consistent orientation of shared edges must therefore be maintained when the elemental contributions are assembled into the global finite-element system.

To examine the tangential behaviour of these functions, consider the edge connecting vertices 1 and 2, with length l_1^e, and let $\mathbf{e}_1$ denote the unit vector oriented from vertex 1 to vertex 2. Along this edge, the barycentric coordinates satisfy [8]

$$\mathbf{e}_1 \cdot \nabla L_1^e = -\frac{1}{l_1^e}, \qquad \mathbf{e}_1 \cdot \nabla L_2^e = \frac{1}{l_1^e}. \tag{3.58}$$

It follows that

$$\mathbf{e}_1 \cdot \mathbf{W}_{12} = \frac{L_1^e + L_2^e}{l_1^e} = \frac{1}{l_1^e}, \tag{3.59}$$

since $L_1^e + L_2^e = 1$ along the edge $(1, 2)$. Thus, $\mathbf{W}_{12}$ has a constant tangential component along its associated edge. On the other two edges, its tangential component vanishes.

Following Jin [8], the basis function associated with edge 1 is therefore defined as

$$\mathbf{N}_1^e = l_1^e \mathbf{W}_{12}, \tag{3.60}$$

so that its tangential component is unity along edge 1 and zero along the other edges. Analogous expressions are obtained for the remaining edges.

The constant tangential behaviour of $\mathbf{W}_{12}$ also implies the standard Nédélec interpolation property

$$\int_{e_k} \mathbf{W}_m \cdot d\mathbf{l} = \delta_{km}. \tag{3.61}$$

Following the formulation presented by Jin [8], the basis function used for the interpolation is obtained by scaling $\mathbf{W}_{ij}$ with the length of the corresponding edge. For example, for the edge connecting vertices 1 and 2,

$$\mathbf{N}_1^e = l_1^e \mathbf{W}_{12}, \tag{3.62}$$

where l_1^e is the length of the edge. Analogous expressions are obtained for the remaining edges of the triangular element.

The electric field within an element is then approximated as

$$\mathbf{E}^e(\mathbf{x}) = \sum_{m=1}^{N_e} \mathbf{N}_m^e(\mathbf{x}) E_m^e, \tag{3.63}$$

where N_e is the number of edges of the element, $\mathbf{N}_m^e$ is the vector basis function associated with edge m, and E_m^e is the corresponding coefficient associated with the tangential component of the electric field.

This expression is the edge-element counterpart of the nodal interpolation

$$u^e(\mathbf{x}) = \sum_{i=1}^{N_n} u_i^e N_i^e(\mathbf{x}), \tag{3.64}$$

introduced in section 3.2.

This construction provides the tangential continuity required for the electromagnetic field across adjacent elements [12, 8].

3.5.2 Weak Formulation for Edge Elements

While the nodal formulation presented previously is based on the scalar TE and TM equations, edge elements can be applied directly to the vector wave equation [8]. Considering the electric-field formulation introduced in eq. (2.20),

$$\nabla \times \left(\mu_r^{-1} \nabla \times \mathbf{E}\right) - k_0^2 \varepsilon_r \mathbf{E} = -ik_0 \eta_0 \mathbf{J}, \tag{3.65}$$

the unknown is now the vector electric field $\mathbf{E}$ rather than a single scalar component.

Following the Galerkin approach [8], the equation is multiplied by a vector-valued test function $\mathbf{W}_i$ and integrated over the element domain Ω^e,

$$\int_{\Omega^e} \mathbf{W}_i \cdot \nabla \times \left(\mu_r^{-1} \nabla \times \mathbf{E}\right) \, d\Omega - k_0^2 \int_{\Omega^e} \mathbf{W}_i \cdot (\varepsilon_r \mathbf{E}) \, d\Omega = -ik_0 \eta_0 \int_{\Omega^e} \mathbf{W}_i \cdot \mathbf{J} \, d\Omega. \tag{3.66}$$

The first term contains the curl of the curl of the electric field. In order to obtain the weak form, the curl integration-by-parts identity is employed [8],

$$\int_{\Omega} \mathbf{A} \cdot (\nabla \times \mathbf{B}) \, d\Omega = \int_{\Omega} (\nabla \times \mathbf{A}) \cdot \mathbf{B} \, d\Omega - \int_{\Gamma} (\mathbf{n} \times \mathbf{A}) \cdot \mathbf{B} \, d\Gamma, \tag{3.67}$$

where $\mathbf{n}$ denotes the outward unit normal to the boundary.

Applying this identity to the first term of eq. (3.66) gives

$$\begin{aligned} &\int_{\Omega^e} (\nabla \times \mathbf{W}_i) \cdot \mu_r^{-1} (\nabla \times \mathbf{E}) \, d\Omega - k_0^2 \int_{\Omega^e} \mathbf{W}_i \cdot (\varepsilon_r \mathbf{E}) \, d\Omega \\ &\quad - \int_{\Gamma^e} (\mathbf{n} \times \mathbf{W}_i) \cdot \mu_r^{-1} (\nabla \times \mathbf{E}) \, d\Gamma = -ik_0 \eta_0 \int_{\Omega^e} \mathbf{W}_i \cdot \mathbf{J} \, d\Omega. \end{aligned} \tag{3.68}$$

As in the nodal formulation, the boundary term accounts for the appropriate boundary conditions and the contributions from internal interfaces are handled during the assembly of the global finite element system. The important difference is that the weak formulation now involves the curl of the vector basis functions rather than the gradient of scalar shape functions.

Using the Nédélec interpolation introduced in eq. (3.63), and choosing the same Nédélec functions as test functions, the Galerkin discretization leads to an elemental system of the form

$$\mathbf{A}^e\mathbf{E}^e = \mathbf{b}^e, \tag{3.69}$$

where the volume contribution to the elemental matrix is

$$\begin{aligned}(A^e_\Omega)_{ij} = &\int_{\Omega^e} (\nabla \times \mathbf{N}^e_i) \cdot \mu_r^{-1} \left(\nabla \times \mathbf{N}^e_j\right)\, d\Omega \\ &- k_0^2 \int_{\Omega^e} \mathbf{N}^e_i \cdot \left(\boldsymbol{\varepsilon}_r \mathbf{N}^e_j\right)\, d\Omega,\end{aligned} \tag{3.70}$$

and the volume source contribution is

$$b^e_{\Omega,i} = -ik_0\eta_0 \int_{\Omega^e} \mathbf{N}^e_i \cdot \mathbf{J}\, d\Omega. \tag{3.71}$$

The boundary integral arising from integration by parts is retained separately and treated according to the boundary condition imposed on Γ^e. Contributions on internal element interfaces are accounted for during global assembly, while contributions on the external boundary enter the discrete system according to the prescribed electromagnetic boundary condition.

4 Developments in the EMWAVE Module

The developments presented in this work were implemented within `EMWAVE`, the electromagnetic module of the Alya simulation code. Before describing these developments, this section briefly introduces the Alya framework, the organization of `EMWAVE`, and the input files used to define an electromagnetic problem.

4.1 Alya and the EMWAVE Module

Alya is a high-performance multiphysics simulation code developed at the Barcelona Supercomputing Center (BSC) for the numerical solution of large-scale scientific and engineering problems [18]. The code follows a modular architecture in which different physical problems are handled by dedicated modules while sharing a common computational infrastructure. This infrastructure provides functionalities including mesh management, parallelization, algebraic solvers, and input/output operations [18].

Alya was designed for distributed-memory high-performance computing architectures, and its scalability has been demonstrated for large multiphysics simulations using up to approximately 10^5 processing cores [18]. Its parallel infrastructure is also employed by `EMWAVE`; the parallel performance of this module is investigated separately on MareNostrum 5 in Chapter 6.

`EMWAVE` provides a finite-element framework for frequency-domain electromagnetic simulations and supports the complex material properties required for the treatment of dielectric and plasma media [9]. The module contains the routines specific to the electromagnetic problem while relying on the Alya kernel for operations such as mesh management, global assembly, algebraic solution, and parallel execution.

4.2 Input Files and Problem Definition

The electromagnetic problem in `EMWAVE` is defined through the module-specific `.max.dat` input file. This file provides the information required by the module to set up and solve a given electromagnetic case, including its physical configuration, numerical settings, requested output quantities, and boundary conditions.

The first part of the `.max.dat` file defines the physical configuration of the problem. Here, the dimensionality of the simulation, the excitation frequency, the wave polarization, and the material properties can be specified. Each material is identified separately, and its permittivity and permeability can be prescribed through their real and imaginary tensor components. This allows both isotropic materials and anisotropic media, such as the cold-plasma dielectric response introduced in section 2.4, to be represented.

The numerical treatment is specified in a separate part of the file. This includes the configuration of the algebraic solver and its convergence parameters. For example, the solver tolerance, maximum number of iterations, and preconditioner can be selected according to the requirements of the simulation.

The quantities written to output files after the electromagnetic problem has been solved are selected in the post-processing part of the file. Depending on the requested output, the real, imaginary, and magnitude components of the available electromagnetic quantities can be written for subsequent visualization and analysis.

Finally, the electromagnetic boundary conditions are associated with the corresponding boundary codes of the computational mesh.

While the `.max.dat` file contains the information specific to `EMWAVE`, the computational domain is defined through the general Alya input files. The `.dom.dat` file specifies the dimensions of the problem and links the files containing the mesh geometry, sets, and boundary-condition information.

The mesh is described in the `.geo.dat` file. In addition to the nodal coordinates and element connectivity, this file associates each finite element with a material identifier. These identifiers provide the connection between the geometrical regions of the mesh and the material properties prescribed in the `.max.dat` file. Boundary information is likewise associated with the corresponding mesh entities and used when the electromagnetic boundary conditions are imposed.

4.3 Computational Workflow of EMWAVE

At the beginning of a simulation, `EMWAVE` initializes the variables and data structures required by the electromagnetic problem and reads the problem definition described in section 4.2.

The electromagnetic problem is then discretized using the nodal or edge-element formulations introduced in Chapter 3. The corresponding finite-element contributions are constructed at element level and assembled into the global algebraic system. The resulting complex-valued system is solved using the algebraic solver infrastructure provided by Alya. After the electromagnetic solution has been obtained, `EMWAVE` performs the

required post-processing operations and writes the selected output variables for subsequent visualization and analysis.

The general computational workflow is summarized in fig. 3. This workflow also provides the context for the developments presented in the following sections, which affect different stages of the `EMWAVE` solution process.

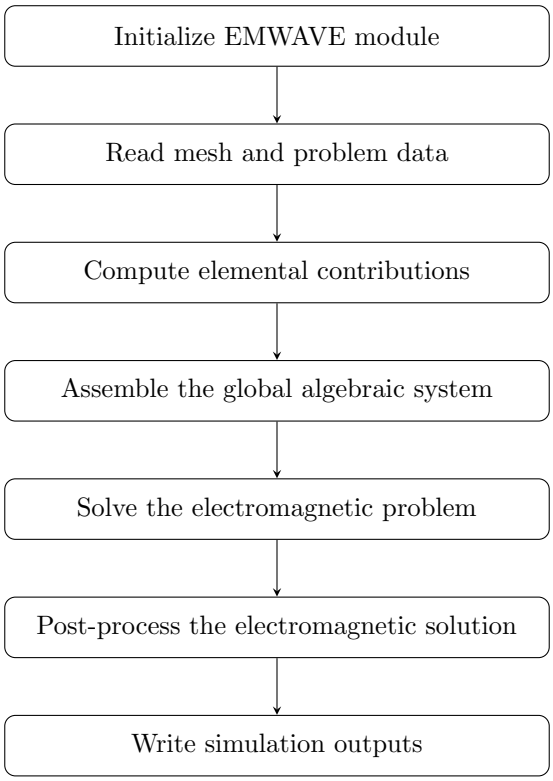

Figure 3: Simplified computational workflow of an `EMWAVE` simulation within Alya.

4.4 Extension to Quadratic and Quadrilateral Elements

The finite-element formulation presented in Chapter 3 allows the solution of the electromagnetic problem using different element topologies and interpolation orders [8]. Prior to this work, however, the nodal formulation of `EMWAVE` had primarily been employed and tested using first-order triangular elements. One of the first objectives of the internship was therefore to extend the range of discretizations effectively employed and verified within the module.

Two additional discretizations were considered: six-node quadratic triangular elements and four-node bilinear quadrilateral elements. Their mathematical formulation, master elements, and corresponding shape functions were introduced in Chapter 3.

It is important to distinguish between the finite-element infrastructure of Alya and its use within the `EMWAVE` module. Alya already provided the general geometrical and finite-element framework required to handle different element topologies and interpolation orders [19]. In particular, the element type determines quantities such as the number of nodes and the number of integration points per element, while the corresponding shape functions, derivatives, quadrature rules, Jacobian transformations, and geometrical information are provided by the common Alya infrastructure.

Therefore, the objective of this development was not to introduce new finite-element definitions or shape functions into Alya itself. Instead, the existing capabilities of the Alya kernel were exploited and their effective use within the nodal electromagnetic formulation of `EMWAVE` was investigated and verified.

The elemental operations performed by `EMWAVE` are already expressed in terms of element-dependent quantities. In particular, the number of nodes of the current element, `pnode`, and the corresponding number of Gauss integration points, `pgaus`, are obtained from the element description. The elemental matrices and right-hand-side vectors can therefore be constructed by looping over the appropriate number of nodes and integration points without requiring a completely independent formulation for every element topology.

This generic structure made it possible to exploit the element definitions already available in Alya and to introduce dedicated `EMWAVE` test cases for quadratic triangular and bilinear quadrilateral discretizations.

As part of this development, the corresponding computational meshes were generated using GiD [23]. This required to become familiar with the mesh-generating workflow in GiD, and in particular, with the generation and export of meshes with the element topologies and interpolation order required for the studied `EMWAVE` test case.

4.5 Reference Cold-Plasma Case

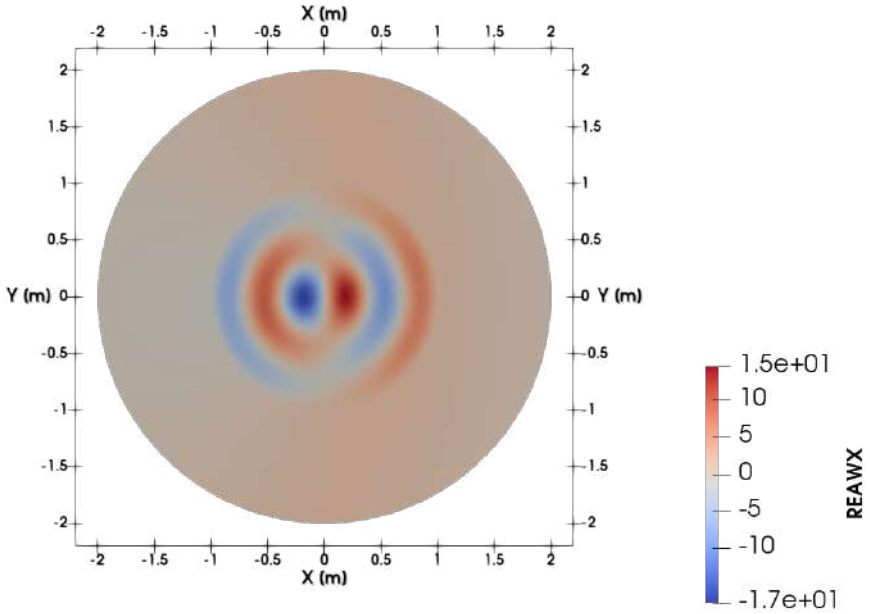

(a) Real component of the electric field.

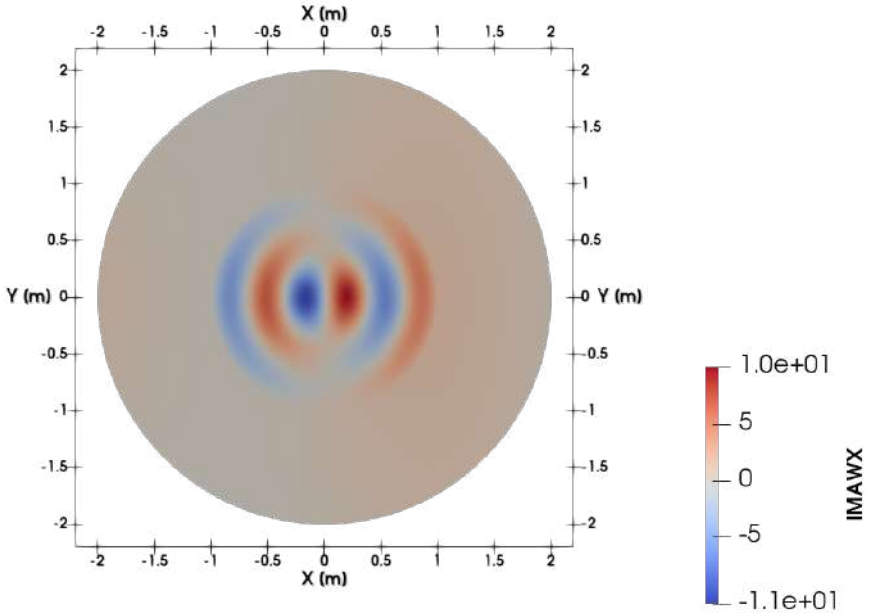

(b) Imaginary component of the electric field.

Figure 4: Real and imaginary components of the electric field for the reference case using first-order triangular finite elements. Extracted from [24].

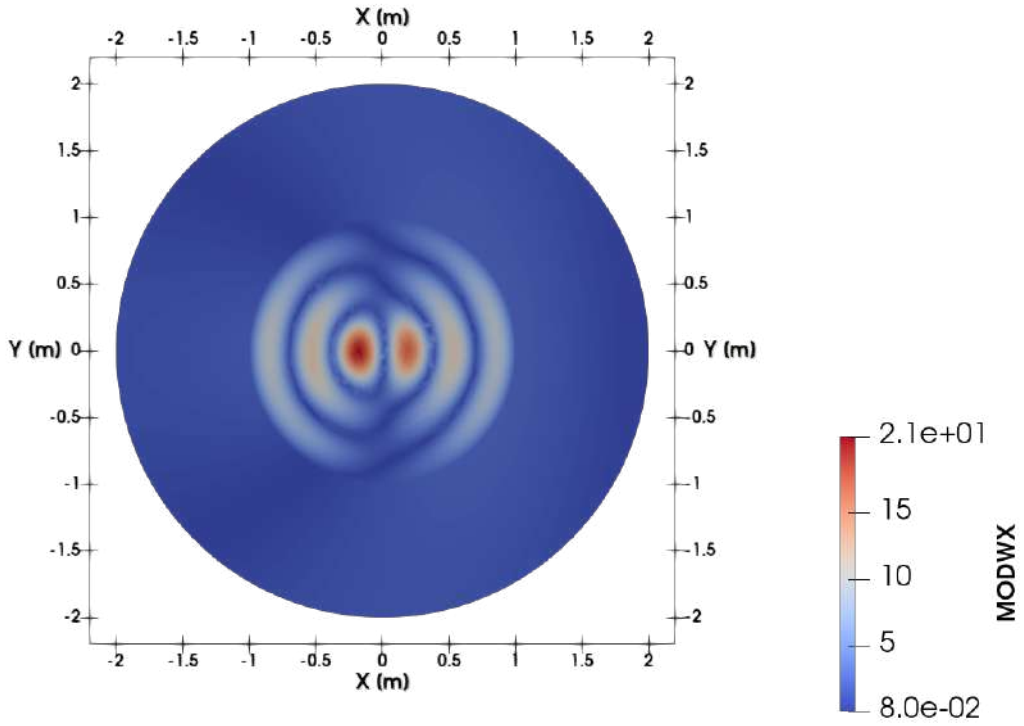

Figure 5: Electric-field magnitude obtained with the reference `TS-CIRC_COLDPLASMA` case using first-order triangular finite elements. The field distribution is shown over the two-dimensional computational domain. Extracted from [24].

The case consists of a two-dimensional circular geometry divided into two concentric regions, with a homogeneous magnetized cold plasma in the inner domain and vacuum in the surrounding annular region. The simulation considers an extraordinary (X) wave at a frequency of 60 MHz, with the plasma response described by the anisotropic cold-plasma permittivity tensor introduced in section 2.4.

The original mesh of the test case contains 2870 first-order triangular elements. For the comparisons presented in this chapter, the spatial resolution was increased using `MULTIPLICATION = 1`, resulting in 11480 first-order triangular elements and 5867 complex-valued degrees of freedom. This refined linear triangular solution is used as the reference for the assessment of the alternative finite-element discretizations introduced in this work.

The real and imaginary parts of the electric field for this test case are shown in figs. 4a and 4b, while its magnitude is shown in Figure 5.

This configuration provides the reference solution used in the following sections to evaluate the quadratic triangular and bilinear quadrilateral discretizations while maintaining the same physical configuration.

4.6 Quadratic Triangular Elements

The first discretization investigated was the six-node quadratic triangular element. As discussed in Chapter 3, this element contains the three vertices of a conventional linear triangle together with three additional nodes located along its edges. The resulting second-order interpolation provides a richer approximation space than that of the three-node linear triangle.

The use of quadratic triangular elements in `EMWAVE` was assessed by introducing the test case

`TS-CIRC_COLDPLASMA_QUADRATIC`.

The case considers the same circular two-dimensional cold-plasma problem previously solved using first-order triangular elements, while employing a quadratic triangular discretization. Maintaining the same physical configuration makes it possible to assess the behaviour of the electromagnetic solver when the interpolation order of the spatial discretization is increased.

The quadratic case was subsequently incorporated into the `EMWAVE` regression test suite. Reference outputs were generated and stored together with the test case, allowing future modifications of the module to be checked for unintended changes in the numerical solution.

The resulting electromagnetic field reproduced the same characteristic spatial structure as the corresponding first-order triangular case, while providing a visibly more detailed representation of the spatial variations of the solution. This behaviour is consistent with the higher-order interpolation provided by the six-node element. A more detailed comparison of the numerical solutions is presented in Chapter 5.

4.7 Bilinear Quadrilateral Elements

The second discretization considered was the four-node bilinear quadrilateral element. In contrast to the quadratic triangular element, the quadrilateral element changes the topology of the mesh while retaining a first-order interpolation along each coordinate of its master element.

The use of this discretization was initially assessed through the test case

`TS-CIRC_DIELECTRIC_QUADRILATERAL`,

which reproduces a circular two-dimensional dielectric problem using a quadrilateral mesh. To provide a comparison with a triangular discretization at a similar mesh resolution, the corresponding first-order triangular case was rerun with `MULTIPLICATION`=1. This increased the number of triangular elements from 2870 to 11480, corresponding to 5867 degrees of freedom. The quadrilateral mesh contains 7933 elements and 8185 degrees of freedom. Although the two discretizations therefore have numbers of elements of the same order, their meshes and numbers of degrees of freedom are not identical.

The corresponding two-dimensional field distributions are shown in Figures 6a and 6b, that show that the same overall electromagnetic field structure is recovered with both meshes.

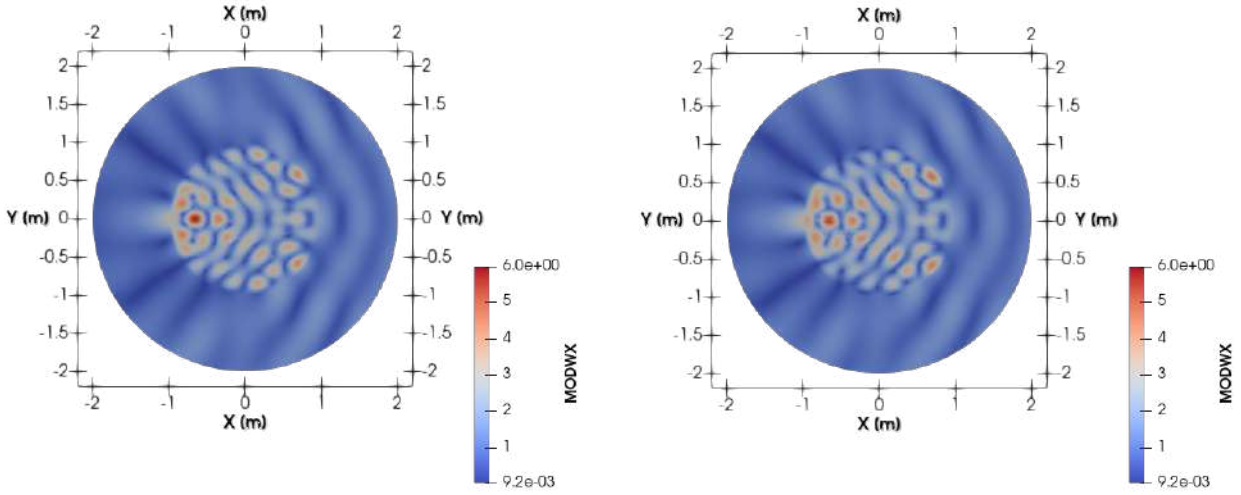

(a) First-order triangular mesh with `MULTIPLICATION`=1 (11480 elements).

(b) Bilinear quadrilateral mesh (7933 elements and 8185 degrees of freedom).

Figure 6: Electric-field magnitude for the circular dielectric case using first-order triangular and bilinear quadrilateral discretizations. Extracted from ParaView [24].

The electric-field magnitude obtained with the two discretizations is compared along the same diameter of the computational domain in Figure 7. The two curves reproduce the same oscillatory structure, including the locations of the main maxima and minima. Differences remain in the amplitudes of some of the extrema, particularly in the left half of the domain, but the overall behaviour of the solution is consistently reproduced by both discretizations.

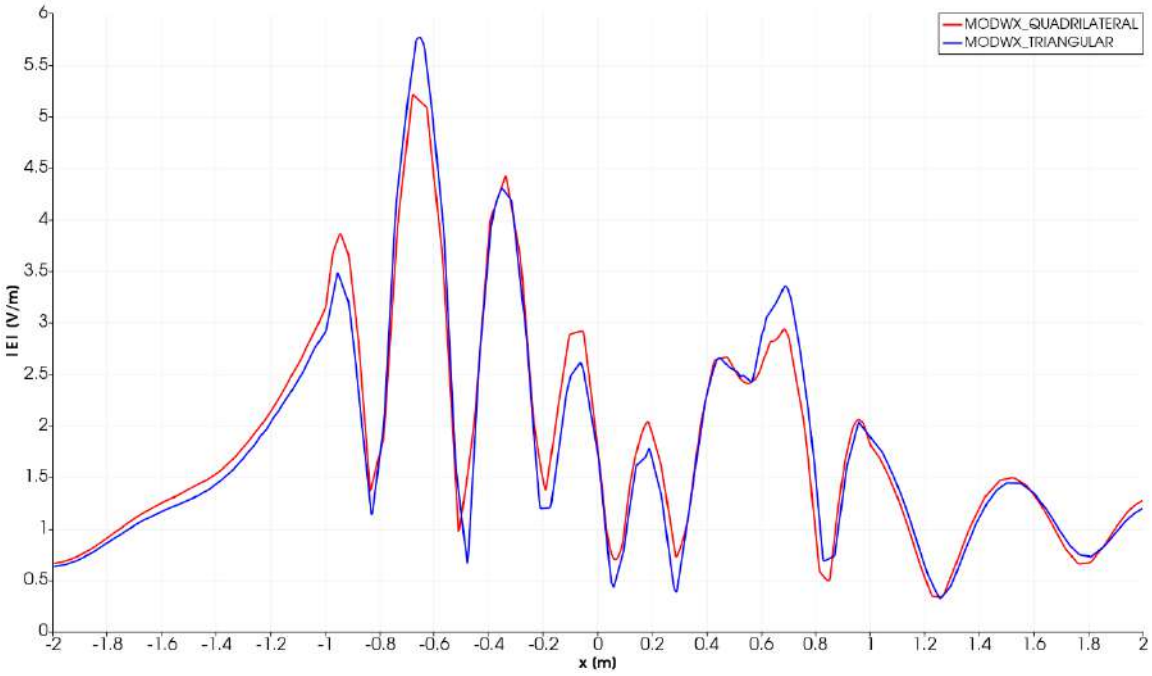

Figure 7: Comparison of the electric-field magnitude along the diameter $y = 0$, from $x = -2$ m to $x = 2$ m, for the first-order triangular and bilinear quadrilateral discretizations of the circular dielectric case. Red corresponds to the quadrilateral case, and blue to the triangular case. Extracted from ParaView [24].

The remaining differences between the two solutions should not be interpreted as arising exclusively from the element topology. Although the refinement of the triangular case provides a number of elements of the same order as that of the quadrilateral mesh, the two meshes are not identical in either their element distribution or their number of degrees of freedom. The comparison therefore verifies that the bilinear quadrilateral discretization reproduces the main features of the electromagnetic solution obtained with the first-order triangular formulation, rather than constituting a formal mesh-convergence comparison.

Since the applications considered in this work are primarily concerned with wave propagation in magnetized plasmas, an additional test was performed to determine whether the quadrilateral discretization could also be employed with the cold-plasma permittivity tensor. For this purpose, a new case,

`TS-CIRC_COLDPLASMA_QUADRILATERAL`,

was constructed from the previously verified quadrilateral configuration, replacing the dielectric material properties with those of the cold-plasma model.

The resulting field distribution showed good qualitative agreement with the corresponding cold-plasma solutions obtained using triangular elements. In particular, the characteristic spatial structure of the electromagnetic field was recovered with the quadrilateral discretization, providing an additional verification that bilinear quadrilateral elements can be employed together with the cold-plasma formulation.

4.8 Comparison of the Discretizations

Figures 8 and 9 compare the magnitude of the electric field obtained for the circular cold-plasma configuration using the quadratic triangular and bilinear quadrilateral elements.

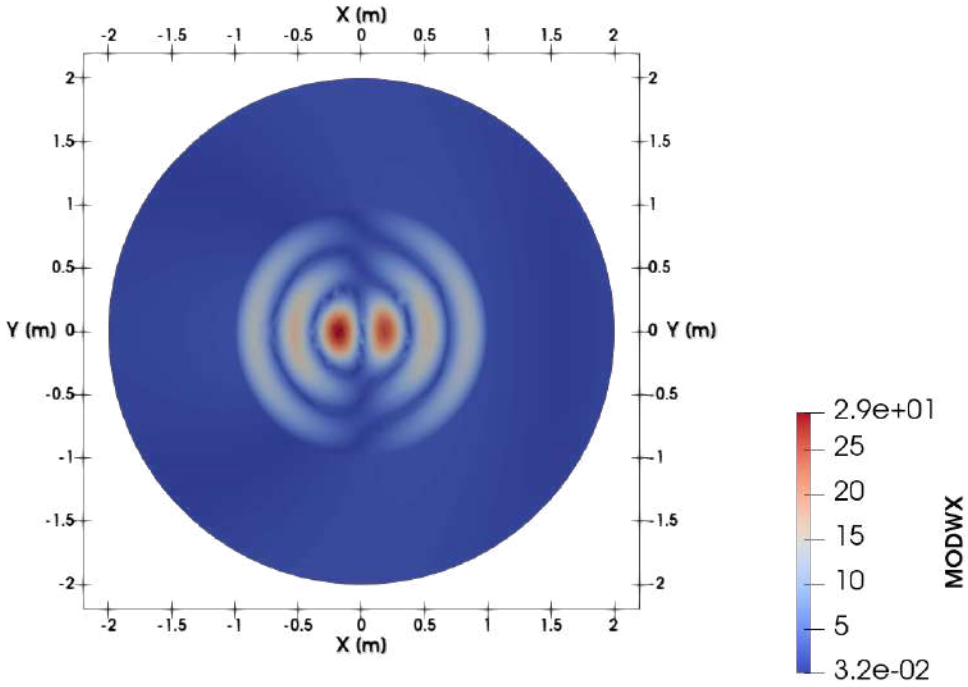

Figure 8: Electric-field magnitude obtained with the second-order triangular discretization. Extracted using ParaView [24].

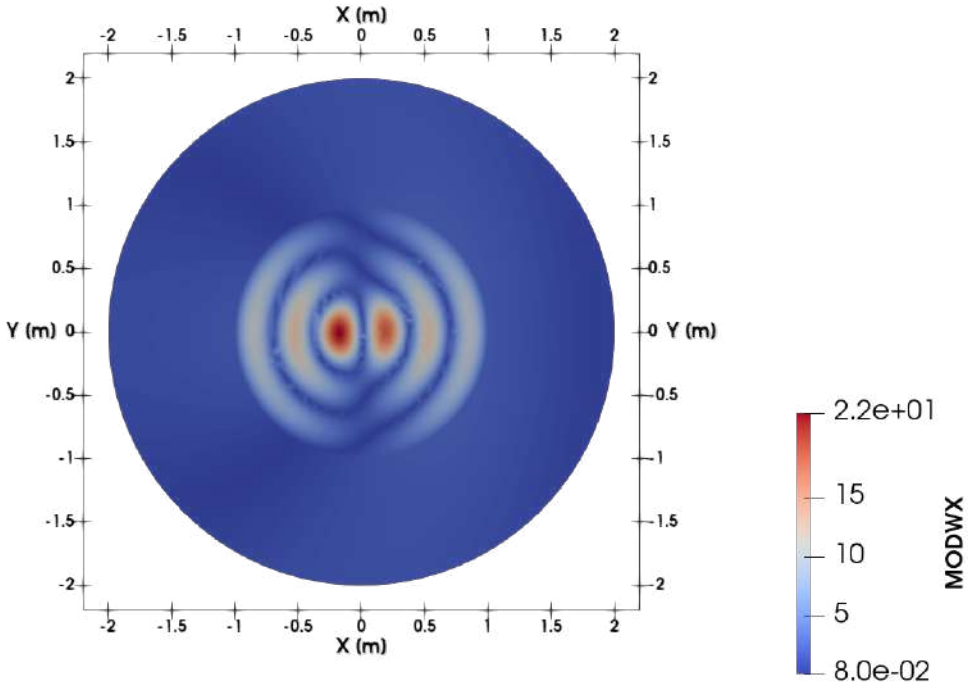

Figure 9: Electric-field magnitude obtained with the bilinear quadrilateral discretization. Extracted using ParaView [24].

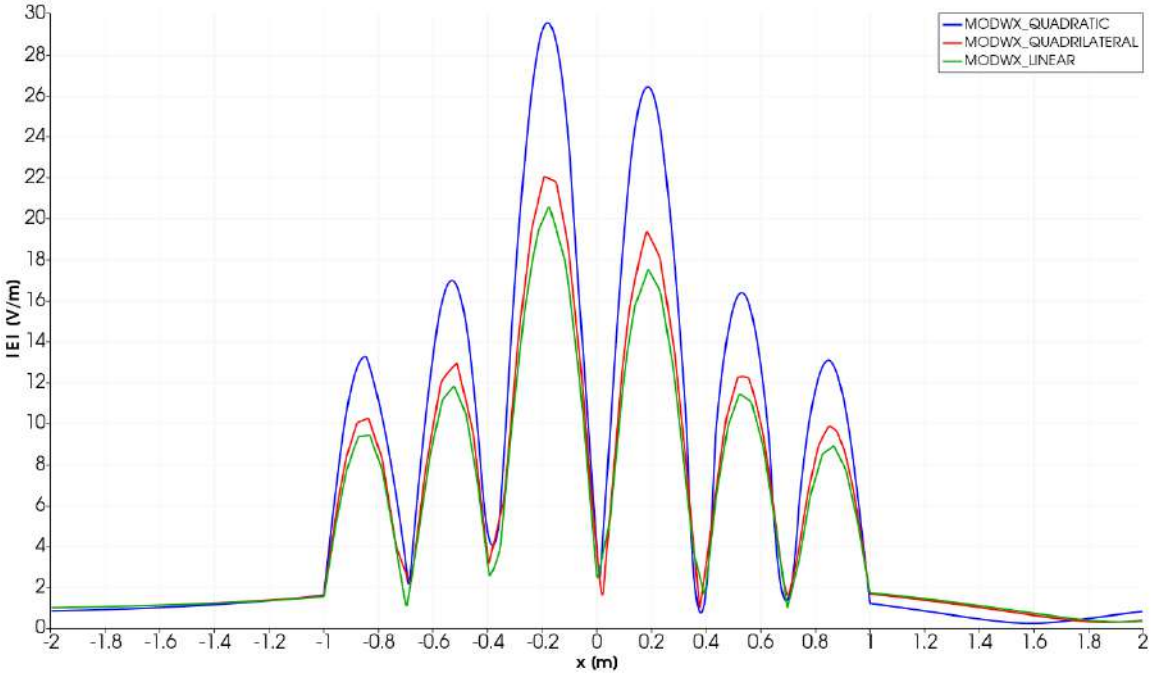

Figure 10: Comparison of the electric-field magnitude along the diameter $y = 0$, from $x = -2$ m to $x = 2$ m, for the first-order triangular, second-order triangular, and bilinear quadrilateral discretizations. Green corresponds to the first-order triangular reference case, red to the quadrilateral case, and blue to the second-order triangular case. Extracted using ParaView [24]

The second-order triangular mesh contains 2788 elements and 5785 complex-valued degrees of freedom, which is very close to the 5867 degrees of freedom of the refined first-order triangular reference case. The bilinear quadrilateral discretization contains 7933 elements and 8185 complex-valued degrees of freedom.

A more direct comparison is provided by the line profiles shown in Figure 10. For all three discretizations, the electric-field magnitude was extracted along the same diameter of the circular domain, extending from $(-2, 0)$ to $(2, 0)$. The first-order triangular and bilinear quadrilateral solutions, represented by the green and red curves, respectively, show a similar spatial behaviour, although visible differences remain along the profile. The second-order triangular solution, shown in blue, follows the same overall field distribution but reaches slightly larger amplitudes, particularly around the central peak, and exhibits a smoother profile.

Since the first and second-order triangular cases have a very similar number of degrees of freedom, the smoother behaviour of the quadratic solution cannot be explained simply by a larger number of unknowns. It is instead consistent with the higher-order interpolation of the quadratic elements, which allows the variation of the field within each element to be represented with greater flexibility. This may also contribute to the slightly larger value obtained around the central maximum. Establishing which solution provides the more accurate estimate of the peak would, however, require a dedicated mesh-convergence analysis.

Overall, the three discretizations reproduce the same characteristic spatial structure of the electromagnetic solution. The agreement obtained between the alternative discretizations therefore supports the correct operation of the implemented second-order triangular and bilinear quadrilateral elements for the anisotropic cold-plasma configuration.

4.9 Implementation of Electrical Conductivity

The theoretical treatment of conductive media and its representation through an effective complex permittivity are presented in section 2.5. The present section therefore focuses on the modifications introduced in `EMWAVE` to incorporate electrical conductivity into the existing material description.

Prior to this development, the electromagnetic material properties handled by `EMWAVE` included relative permittivity and permeability, whereas electrical conductivity could not be prescribed directly through the physical-data input. To introduce this additional material property, the data structures of the module were first extended in `def_maxwel.f90` with the variable `sigma_material`. This variable stores the electrical conductivity associated with each material defined in the computational domain.

The input-reading procedure in `max_reaphy.f90` was subsequently extended with the keyword `SIGMA`, allowing the conductivity to be specified together with the other electromagnetic material properties. The value assigned to each material is stored in the corresponding entry of `sigma_material` and is then used during the finite-element calculations to construct the effective complex permittivity introduced in section 2.5. The implementation was incorporated into both the nodal and edge-element formulations of `EMWAVE`.

Since Alya is designed for parallel execution [19], the new material property also had to be incorporated into the data communication stage. During the verification of the implementation, it was found that, although the conductivity values were correctly read from the input data, they were not initially communicated to all MPI processes. As a consequence, different processes could have inconsistent conductivity information during a parallel simulation.

To correct this behaviour, `sigma_material` was added to the module-specific data communication performed in `max_sendat.f90`. The conductivity values associated with the different materials are therefore distributed consistently among the MPI processes before the electromagnetic system is assembled.

With these modifications, conductivity is integrated throughout the material-data workflow of `EMWAVE`. For each material, the value of σ is read from the physical-data input, stored in `sigma_material`, communicated to the MPI processes, and used during the finite-element assembly. The same material property is also used during post-processing for the computation of the conduction current density $\mathbf{J}$, following eq. (2.55).

A dedicated test case, `TS-RECT_WAVEGUIDE_CONDUCTIVITY`, was introduced to exercise the new implementation. The case consists of a rectangular waveguide containing a material with finite prescribed conductivity and was used to verify the complete conductivity workflow, from the input data to the electromagnetic solution. Reference outputs were subsequently included in the `EMWAVE` regression-test framework, allowing future modifications of the module to be checked for unintended changes in conductive-material simulations.

The quantitative validation of the conductivity implementation, including comparisons of the resulting electromagnetic fields with the ERMES 20.0 solver, is presented in Chapter 5.

4.10 Implementation of Additional Electromagnetic Fields

Another extension introduced in the `EMWAVE` module concerned the computation and post-processing of additional electromagnetic quantities. While the electric field $\mathbf{E}$ constitutes the primary unknown of the vector electromagnetic formulation, the magnetic flux density $\mathbf{B}$, the magnetic field strength $\mathbf{H}$, and the current density $\mathbf{J}$ are also of direct physical interest. These quantities are not introduced as additional unknowns in the finite-element system, but are recovered from the computed electric-field solution using Maxwell's equations and the corresponding constitutive relations [8].

The implementation described in this section concerns the edge-element formulation of `EMWAVE`; the additional fields were not implemented for the nodal formulation.

The magnetic flux density is obtained from Faraday's law in its time-harmonic form, introduced in Eq. (2.9). Solving for $\mathbf{B}$ gives

$$\mathbf{B} = \frac{1}{i\omega}\nabla \times \mathbf{E}. \tag{4.1}$$

Once $\mathbf{B}$ has been reconstructed, the magnetic field strength follows from the constitutive relation in Eq. (2.6). For the isotropic magnetic permeability considered in the present implementation,

$$\mathbf{H} = \frac{\mathbf{B}}{\mu_0 \mu_r}. \tag{4.2}$$

The current density considered in this implementation is the conduction current density. For an isotropic conductive medium, it follows directly from Ohm's law, introduced in eq. (2.55).

The implementation required extending the existing post-processing workflow so that the newly computed quantities could be reconstructed and made available as simulation outputs. The main stages of this implementation are summarized in Figure 11.

The required data structures were first introduced in `def_maxwel.f90`. Since `EMWAVE` solves the time-harmonic electromagnetic problem using complex-valued fields, separate arrays were defined for the real and imaginary components of each new quantity. Two representations were introduced for each additional field. At the element level, the real and imaginary components are stored separately in `reahx_gp_max` and `imahx_gp_max` for $\mathbf{H}$, `reabx_gp_max` and `imabx_gp_max` for $\mathbf{B}$, and `reajx_gp_max` and `imajx_gp_max` for $\mathbf{J}$. The corresponding nodal representations are stored in `reahx_node` and `imahx_node`, `reabx_node` and `imabx_node`, and `reajx_node` and `imajx_node`, respectively.

Figure 11: Overview of the implementation steps required to compute and provide the fields **B**, **H**, and **J** in `EMWAVE`. The pipeline spans from the definition of the new variables to their availability as post-processing outputs.

Memory allocation for these new variables was incorporated in `max_memall.f90`. The nodal arrays are allocated according to the spatial dimension and the number of mesh nodes, whereas the element-based arrays are allocated separately for each element according to the spatial dimension and the number of evaluation points required for that element. This allows the new fields to follow the existing element-dependent data structures used by the edge-element formulation.

The actual reconstruction and computation of the fields is performed in `max_nodedge.f90`. For each element, the global edge indices and their orientations are first retrieved. The Nédélec basis functions introduced in section 3.5 and their curls are then evaluated, allowing both the electric field and its curl to be reconstructed from the edge degrees of freedom. The orientation signs associated with the edges are explicitly taken into account during this operation.

The reconstructed quantities are then used to compute **H**, **B**, and **J** according to the electromagnetic relations introduced in Chapter 2. At code level, the core of this operation can be summarized by the following instructions:

```
hmax_nod = matmul(inv_permea(1:ndime,1:ndime), curlE_nod) / iw_max
bmax_nod = curlE_nod / iw_max
jmax_nod = cmplx(sigma_material(pmate),0.0_rp,kind=rp) * emwax_nod
```

Here, `emwax_nod` and `curlE_nod` denote the reconstructed electric field and its curl, respectively.

The variable `ndime` denotes the number of spatial dimensions of the problem, while `pmate` identifies the material associated with the current element and is used to access its corresponding conductivity through `sigma_material(pmate)`. The intrinsic Fortran function `matmul` performs the matrix–vector product between the inverse permeability tensor and the curl of the electric field, whereas `cmplx` constructs the complex-valued conductivity from its real material value and a zero imaginary component.

In particular, the computation of **H** applies the inverse permeability tensor `inv_permea`, following eq. (4.2), while the calculation of **J** directly uses the material-dependent conductivity stored in `sigma_material`, as in eq. (2.55). The conductivity implementation is therefore connected directly to the post-processing calculation of the conduction current density.

Because the electromagnetic quantities are complex, the real and imaginary components of **E**, **H**, **B**, and **J** are stored separately after their evaluation. The element-based values are subsequently transferred to the mesh nodes. Since a node can belong to several neighbouring elements, the contributions from all elements sharing that node are accumulated. The number of contributing elements is tracked and the accumulated field values are divided by this number, providing an averaged nodal representation of each electromagnetic quantity.

The reconstruction of the field quantities at the mesh nodes also accounts for distributed-memory parallel execution. Before the averaging operation, the accumulated real and imaginary field components, together with

the number of element contributions associated with each node, are communicated between MPI processes. This ensures a consistent representation at nodes shared by different subdomains.

Finally, the new quantities were incorporated into the standard `EMWAVE` post-processing system. The corresponding variables were registered in `max_inivar.f90` and defined as available post-processing quantities in `max_inipos.f90`. The output procedure in `max_outvar_edge.f90` was extended to provide the real and imaginary components of the new fields as well as their magnitudes. In particular, the scalar outputs `MODHX`, `MODBX`, and `MODJX` provide the magnitudes of $\mathbf{H}$, $\mathbf{B}$, and $\mathbf{J}$, respectively, allowing the newly implemented quantities to be directly exported and visualized in ParaView [24].

Together with the developments presented in the preceding sections, these modifications extend the discretization, material modelling, and post-processing capabilities of `EMWAVE`. Chapter 5 assesses the implemented developments through numerical verification and independent validation against ERMES 20.0.

5 Verification and Validation

The developments introduced in the `EMWAVE` module require appropriate verification and validation in order to assess the reliability of the resulting electromagnetic solutions. This is particularly important when introducing new finite-element discretizations, material properties and post-processing quantities, since each modification affects different stages of the numerical solution procedure.

In this work, verification and validation were performed through two complementary approaches. Validation was carried out by comparing selected `EMWAVE` solutions with results obtained using ERMES 20.0 [20], an independent finite-element solver for computational electromagnetics in the frequency domain. Equivalent electromagnetic configurations were reproduced in both codes and the resulting field distributions were compared quantitatively.

The validation procedure was based on a rectangular-waveguide benchmark. Preliminary comparisons were first performed using the original waveguide configuration and a dielectric variant, providing reference checks between `EMWAVE` and ERMES 20.0 before introducing the developments carried out during this work. The validation presented in this chapter focuses on the conductive rectangular waveguide, since this configuration directly assesses the implementation of electrical conductivity described in Chapter 4.

For the conductive case, the comparison was further extended beyond the electric field to the magnetic flux density $\mathbf{B}$, magnetic field $\mathbf{H}$ and current density $\mathbf{J}$. This provides an independent assessment of both the conductivity implementation and the additional electromagnetic field outputs introduced in Chapter 4.

For each comparison, equivalent geometrical and physical parameters were prescribed in `EMWAVE` and ERMES 20.0. The electromagnetic field magnitudes were extracted along the same spatial path through the computational domain. The `EMWAVE` profiles were exported from ParaView and compared with the corresponding ERMES 20.0 profiles as a function of position. This procedure allows both the spatial structure and the magnitude of the electromagnetic solution to be assessed.

Verification was additionally performed for the new quadratic triangular and quadrilateral finite-element configurations by comparing their solutions with reference cases based on the previously available linear triangular discretization.

The objective of this chapter is therefore to assess the numerical behaviour of the developments introduced in `EMWAVE` and quantify their agreement with independent reference solutions.

5.1 Conductive Rectangular Waveguide

The main validation against ERMES 20.0 was performed using the `TS-RECT_WAVEGUIDE_CONDUCTIVITY` test case. The case consists of a three-dimensional rectangular waveguide containing a dielectric material with relative permittivity

$$\varepsilon_r = 10, \tag{5.1}$$

and excited at a frequency

$$f = 60 \text{ MHz}. \tag{5.2}$$

A finite electrical conductivity σ is assigned to the material. Its contribution is incorporated into the electromagnetic material model through the effective complex relative permittivity defined in **??**. Since the material considered in this benchmark is isotropic, the general tensorial formulation reduces to the corresponding scalar case.

An example of the three-dimensional electric-field distribution obtained with `EMWAVE` is shown in Fig. 12.

Equivalent material properties and electromagnetic parameters were prescribed in ERMES 20.0. To perform a quantitative comparison, the field magnitudes obtained with both solvers were extracted along a diagonal line connecting two diametrically opposite vertices of the computational domain. The resulting one-dimensional profiles were then compared as a function of the position along this line.

Several conductivity values were considered during the verification of the implementation. The results presented here include $\sigma = 2$ S/m and $\sigma = 10.3$ S/m. The second value was also selected as an additional check that the new material property could correctly handle arbitrary real-valued conductivities, rather than only integer values.

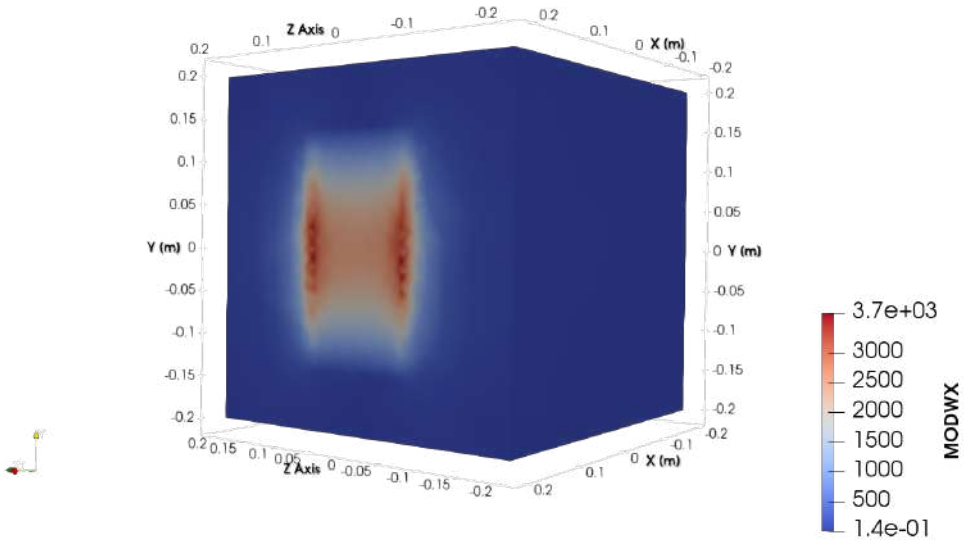

Figure 12: Electric-field magnitude obtained with `EMWAVE` for the three-dimensional conductive waveguide test case, visualized in ParaView [24].

5.2 Numerical Comparison against ERMES

The electromagnetic fields computed with `EMWAVE` were compared against the corresponding ERMES solutions along the selected diagonal of the rectangular waveguide. The comparison was first examined through the field profiles and subsequently quantified using normalized error measures evaluated at the ERMES sampling locations.

The electric field provides the most direct comparison of the electromagnetic solution obtained with the two solvers. Figure 13 shows the magnitude of $\mathbf{E}$ along the selected diagonal for $\sigma = 2$ S/m and $\sigma = 10.3$ S/m.

A very close agreement is obtained between `EMWAVE` and ERMES for both conductivity values. The two solvers reproduce the same spatial variation of the electric field, including the positions of the principal maxima and minima, while the corresponding amplitudes are nearly coincident over the complete extracted profile.

The effect of increasing the conductivity is also clearly visible. For $\sigma = 2$ S/m, the maximum electric-field magnitude is of the order of 90 V/m, whereas for $\sigma = 10.3$ S/m it decreases to approximately 4.5 V/m. Despite this substantial modification of the electromagnetic response, the agreement between the two solvers is maintained. This result provides a particularly useful validation of the conductivity implementation, since `EMWAVE` reproduces not only the field distribution for a single prescribed value of σ, but also the strong change in field amplitude produced by a significantly different conductivity.

The comparison was subsequently extended to the additional electromagnetic fields implemented in `EMWAVE`. As shown in Figure 14, the two solvers reproduce very similar profiles for $|\mathbf{B}|$, $|\mathbf{H}|$, and $|\mathbf{J}|$. For $\mu_r = 1$, the identical spatial dependence of $|\mathbf{B}|$ and $|\mathbf{H}|$ is consistent with $\mathbf{B} = \mu_0\mathbf{H}$. Similarly, since the conductivity is uniform, the current density follows $\mathbf{J} = \sigma\mathbf{E}$ and therefore retains the spatial dependence of the electric field.

Finally, the comparison was repeated after simultaneously changing the conductivity and relative permeability. Figure 15 shows that the agreement is maintained for all four electromagnetic quantities for $\sigma = 10.3$ S/m and $\mu_r = 3$. In particular, this case verifies the computation of $\mathbf{H}$ when the relative permeability differs from unity, while the current-density profile remains consistent with the implemented relation $\mathbf{J} = \sigma\mathbf{E}$.

To quantify the agreement observed in the field profiles, the comparison was complemented by two normalized error measures. Three material configurations were considered,

$$(\sigma, \mu_r) = (2, 1), \qquad (\sigma, \mu_r) = (10.3, 1), \qquad (\sigma, \mu_r) = (10.3, 3). \tag{5.3}$$

This sequence allows the influence of the two material parameters to be examined separately. The transition from $(2, 1)$ to $(10.3, 1)$ modifies only the electrical conductivity, whereas the transition from $(10.3, 1)$ to $(10.3, 3)$ modifies only the relative magnetic permeability.

For each field and material configuration, the comparison was performed at the $N = 24$ original ERMES 20.0 sampling locations along the selected waveguide diagonal. Since the Alya profiles contain a larger number of

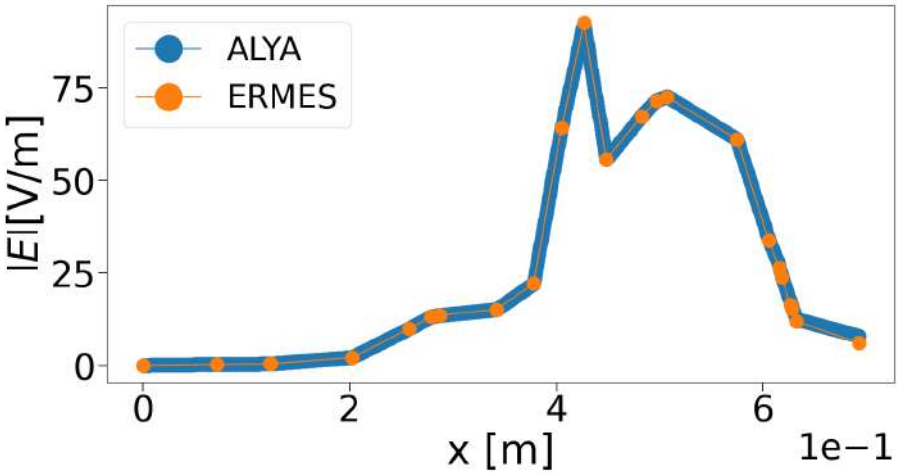

(a) $\sigma = 2$ S/m.

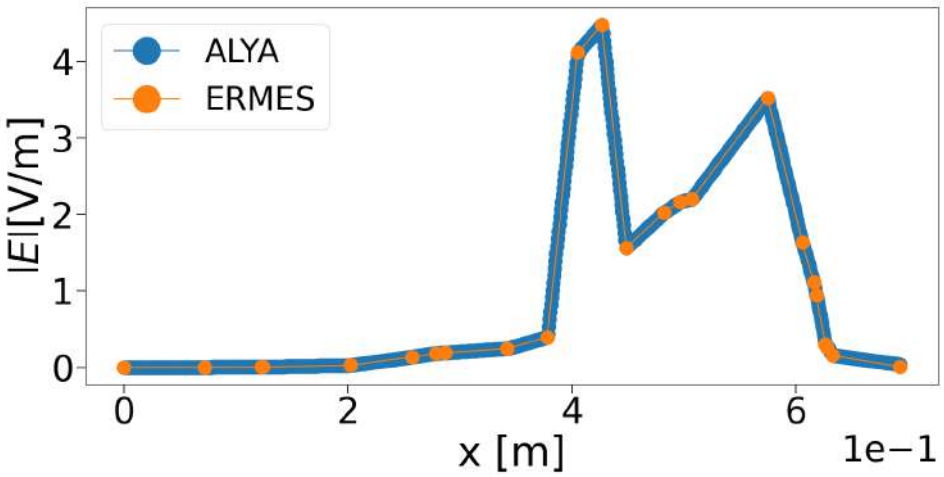

(b) $\sigma = 10.3$ S/m.

Figure 13: Comparison between `EMWAVE` and ERMES of the electric-field magnitude along the selected diagonal of the conductive rectangular waveguide for two values of the electrical conductivity.

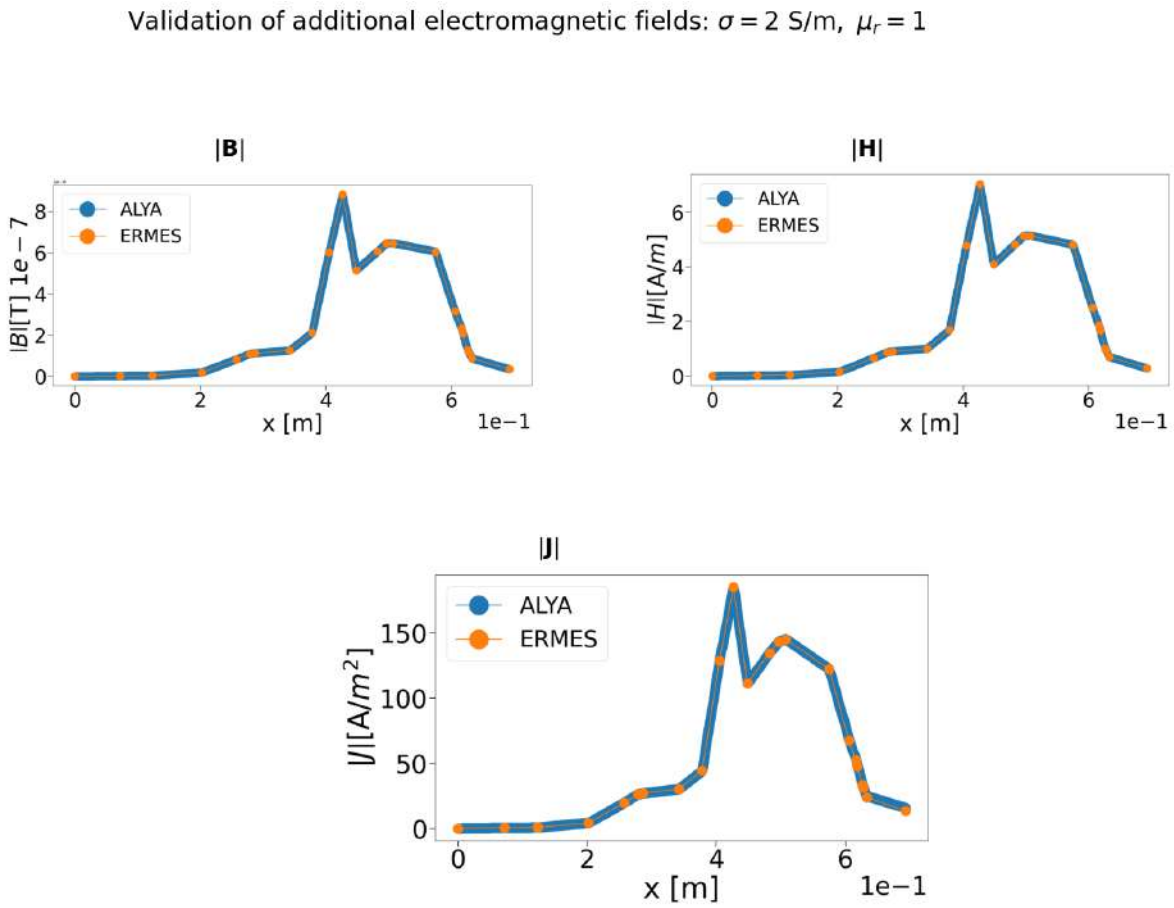

Figure 14: Comparison between `EMWAVE` and ERMES for the magnetic flux density $|\mathbf{B}|$, magnetic field $|\mathbf{H}|$, and current density $|\mathbf{J}|$ in the conductive rectangular waveguide with $\sigma = 2$ S/m and $\mu_r = 1$.

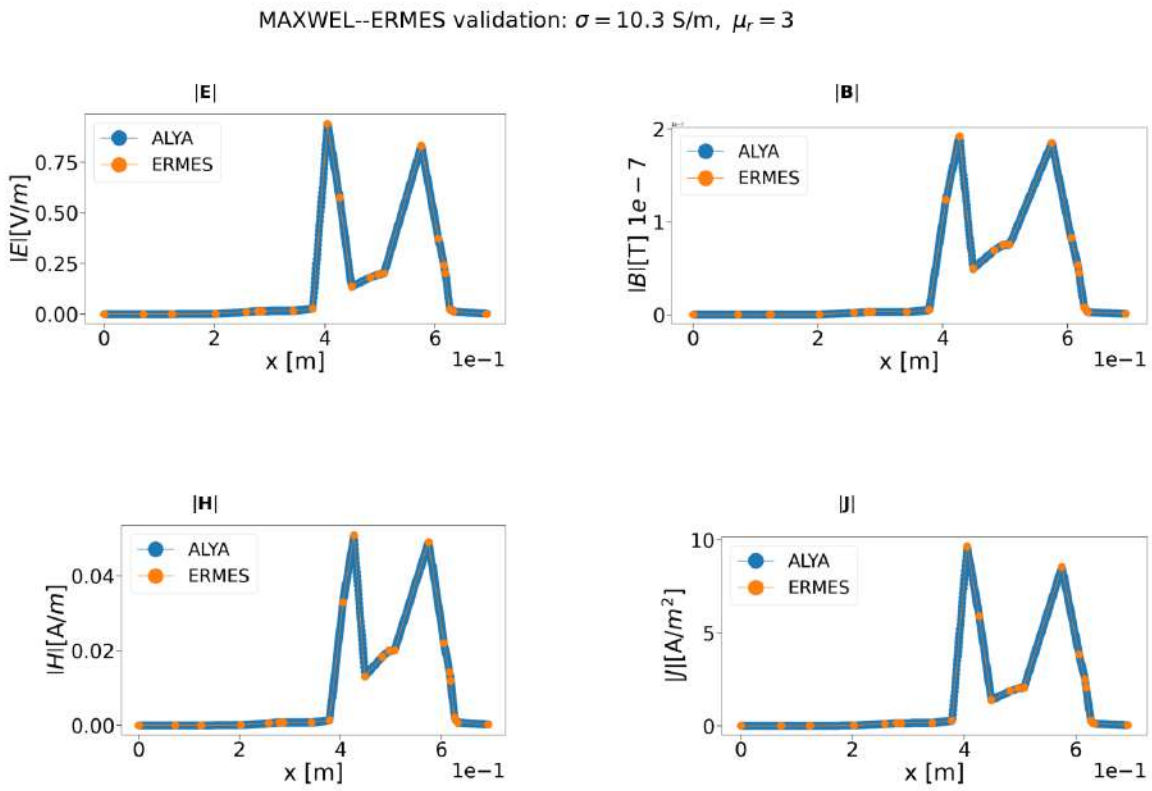

Figure 15: Comparison between `EMWAVE` and ERMES 20.0 for $|\mathbf{E}|$, $|\mathbf{B}|$, $|\mathbf{H}|$, and $|\mathbf{J}|$ with $\sigma = 10.3$ S/m and $\mu_r = 3$.

sampling points, the Alya solution was linearly interpolated at the ERMES 20.0 coordinates before computing the errors.

Two complementary error measures were considered. The first is the root-mean-square error normalized by the maximum absolute amplitude of the ERMES solution,

$$\mathrm{NRMSE} = \frac{\sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(y_i^{\mathrm{Alya}} - y_i^{\mathrm{ERMES}}\right)^2}}{\max_i \left|y_i^{\mathrm{ERMES}}\right|} \times 100. \tag{5.4}$$

The NRMSE provides a global measure of the discrepancy between the two field profiles relative to the maximum amplitude of the reference ERMES 20.0 solution.

The second quantity identifies the largest local discrepancy among the sampled positions,

$$D_{\max} = \frac{\max_i \left|y_i^{\mathrm{Alya}} - y_i^{\mathrm{ERMES}}\right|}{\max_i \left|y_i^{\mathrm{ERMES}}\right|} \times 100. \tag{5.5}$$

Using the maximum ERMES 20.0 amplitude for the normalization, rather than a pointwise relative error, avoids artificially large relative errors in regions where the electromagnetic field approaches zero.

The complete quantitative results for the three material configurations and the four electromagnetic quantities are reported in Table 1.

Table 1: Quantitative comparison between Alya/`EMWAVE` and ERMES 20.0 for the three material configurations considered in the conductive rectangular-waveguide benchmark.

σ [S/m]	μ_r	Field	NRMSE [%]	$D_{\max}$ [%]
2.0	1	$\lvert\mathbf{E}\rvert$	0.3173	0.7738
2.0	1	$\lvert\mathbf{B}\rvert$	0.1123	0.3441
2.0	1	$\lvert\mathbf{H}\rvert$	0.1123	0.3441
2.0	1	$\lvert\mathbf{J}\rvert$	0.2610	0.6248
10.3	1	$\lvert\mathbf{E}\rvert$	0.2266	0.5710
10.3	1	$\lvert\mathbf{B}\rvert$	0.1328	0.4919
10.3	1	$\lvert\mathbf{H}\rvert$	0.1327	0.4915
10.3	1	$\lvert\mathbf{J}\rvert$	0.1718	0.5139
10.3	3	$\lvert\mathbf{E}\rvert$	0.2869	0.9224
10.3	3	$\lvert\mathbf{B}\rvert$	0.1688	0.5487
10.3	3	$\lvert\mathbf{H}\rvert$	0.1688	0.5488
10.3	3	$\lvert\mathbf{J}\rvert$	0.2459	0.9224

The results confirm that the excellent agreement observed visually is maintained over all three material configurations. Increasing the conductivity from 2 to 10.3 S/m while keeping $\mu_r = 1$ does not produce any deterioration of the agreement. For this particular comparison, the NRMSE actually decreases for all four electromagnetic quantities. The largest NRMSE decreases from 0.3173% for $|\mathbf{E}|$ in the $\sigma = 2$ S/m case to 0.2266% when $\sigma = 10.3$ S/m.

The third configuration provides an additional test in which the relative magnetic permeability is increased from $\mu_r = 1$ to $\mu_r = 3$, while the conductivity remains fixed at 10.3 S/m. Slightly larger discrepancies are obtained for some of the fields, but they remain very small. In this configuration, the largest NRMSE is 0.2869%, while the largest value of $D_{\max}$ is 0.9224%.

Considering all twelve comparisons together, the largest NRMSE is

$$\mathrm{NRMSE}_{\max} = 0.3173\%, \tag{5.6}$$

obtained for $|\mathbf{E}|$ at $\sigma = 2$ S/m and $\mu_r = 1$. The largest local normalized deviation is

$$D_{\max,\mathrm{overall}} = 0.9224\%, \tag{5.7}$$

observed for $|\mathbf{E}|$, in the $\sigma = 10.3$ S/m, $\mu_r = 3$ configuration. Therefore, for every field and material configuration considered, the global RMS discrepancy between Alya and ERMES 20.0 remains below approximately

0.32% of the maximum ERMES field amplitude, while even the largest local discrepancy among all sampled positions remains below 1%.

An additional consistency check is provided by the magnetic quantities. Since the permeability is spatially uniform and $\mathbf{B} = \mu\mathbf{H}$, the two fields differ only by a constant factor and their normalized errors are therefore expected to be essentially identical. This behaviour is observed for all three material configurations, providing an additional consistency check of the computed magnetic fields.

Overall, the quantitative results demonstrate that the agreement between Alya/`EMWAVE` and ERMES 20.0 is not restricted to a single set of material parameters. It remains excellent when the electrical conductivity is increased and when a non-unit relative magnetic permeability is introduced. Together with the visual superposition of the field profiles these results provide a robust validation of the newly implemented **B**, **H**, and **J** outputs in `EMWAVE`.

6 Scalability Analysis on MareNostrum 5

One of the main objectives of Alya is the simulation of large-scale multiphysics problems on high-performance computing architectures [19]. In this context, the computational performance of the `EMWAVE` module is an important aspect to consider, particularly for electromagnetic simulations involving large meshes and complex plasma configurations.

During this work, a series of scalability tests was performed on MareNostrum 5 (MN5), the supercomputer hosted at the Barcelona Supercomputing Center [21]. The General Purpose Partition (GPP), used for the simulations presented in this work, comprises 6480 compute nodes and a total of 725760 CPU cores. Each compute node provides 112 CPU cores, corresponding to two 56-core Intel processors.

In the Alya parallel execution model used in this work, one MPI process acts as the master, while the remaining processes operate as workers. For the scalability tests presented here, one MPI process was assigned to each CPU core, so that the number of worker processes corresponds directly to the number of working cores. The master process coordinates the execution, whereas the computational domain is partitioned into subdomains and distributed among the worker processes. Each worker performs the finite-element operations associated with its local portion of the mesh, with the required information exchanged between neighbouring subdomains through MPI communication. Therefore, up to 111 working cores can be employed within a single compute node in the configuration considered here, while simulations using more than 111 workers require the allocation of additional nodes.

The objective of these tests was to characterize the parallel performance of `EMWAVE` when increasing the computational resources assigned to a fixed electromagnetic problem.

The analysis presented in this chapter focuses on strong scaling. In a strong scaling experiment, the size of the numerical problem is kept constant while the number of parallel processes is increased. Ideally, increasing the number of processes by a given factor should reduce the execution time by the same factor. In practice, communication, synchronization, load imbalance and other parallel overheads prevent this ideal behaviour, particularly when the amount of computational work assigned to each process becomes small [22].

The tests presented here were performed using MPI parallelism and different numbers of computational resources on MN5. The execution time, speedup and parallel efficiency were used to evaluate the scaling behaviour of the module.

6.1 Strong Scaling Methodology

Let $T(N)$ denote the execution time obtained using N worker processes. The strong-scaling speedup relative to a reference execution using N_{ref} processes is defined as [22]

$$S(N) = \frac{T(N_{\mathrm{ref}})}{T(N)}. \tag{6.1}$$

For ideal strong scaling, the execution time decreases inversely with the number of processes. The corresponding ideal speedup is therefore

$$S_{\mathrm{ideal}}(N) = \frac{N}{N_{\mathrm{ref}}}. \tag{6.2}$$

A second useful quantity is the parallel efficiency, which measures the fraction of the ideal speedup achieved by a given parallel configuration. It is defined as

$$\eta(N) = \frac{S(N)}{N/N_{\mathrm{ref}}} = \frac{N_{\mathrm{ref}} T(N_{\mathrm{ref}})}{N T(N)}. \tag{6.3}$$

With this definition, an ideally scaling application has $\eta = 1$, while decreasing values of η indicate that the additional computational resources are providing progressively smaller reductions in execution time.

In addition to these quantities, the wall-clock execution time was considered directly. This is particularly useful for identifying changes in scaling behaviour associated with different computational configurations, such as the transition from computations contained within a single computing node to executions distributed across multiple nodes.

6.2 Scaling Results

The strong-scaling behaviour was first evaluated by considering the execution time as a function of the number of MPI processes. As expected, the execution time generally decreases as additional computational resources are assigned to the problem. However, the reduction is not proportional to the increase in the number of processes, indicating a progressive departure from ideal strong-scaling behaviour.

The corresponding speedup was compared with the ideal scaling curve defined in eq. (6.2). The ideal curve represents the theoretical case in which doubling the number of processes results in a factor-of-two reduction in execution time.

6.2.1 TS-CIRC_DIELECTRIC_ISOTROPIC

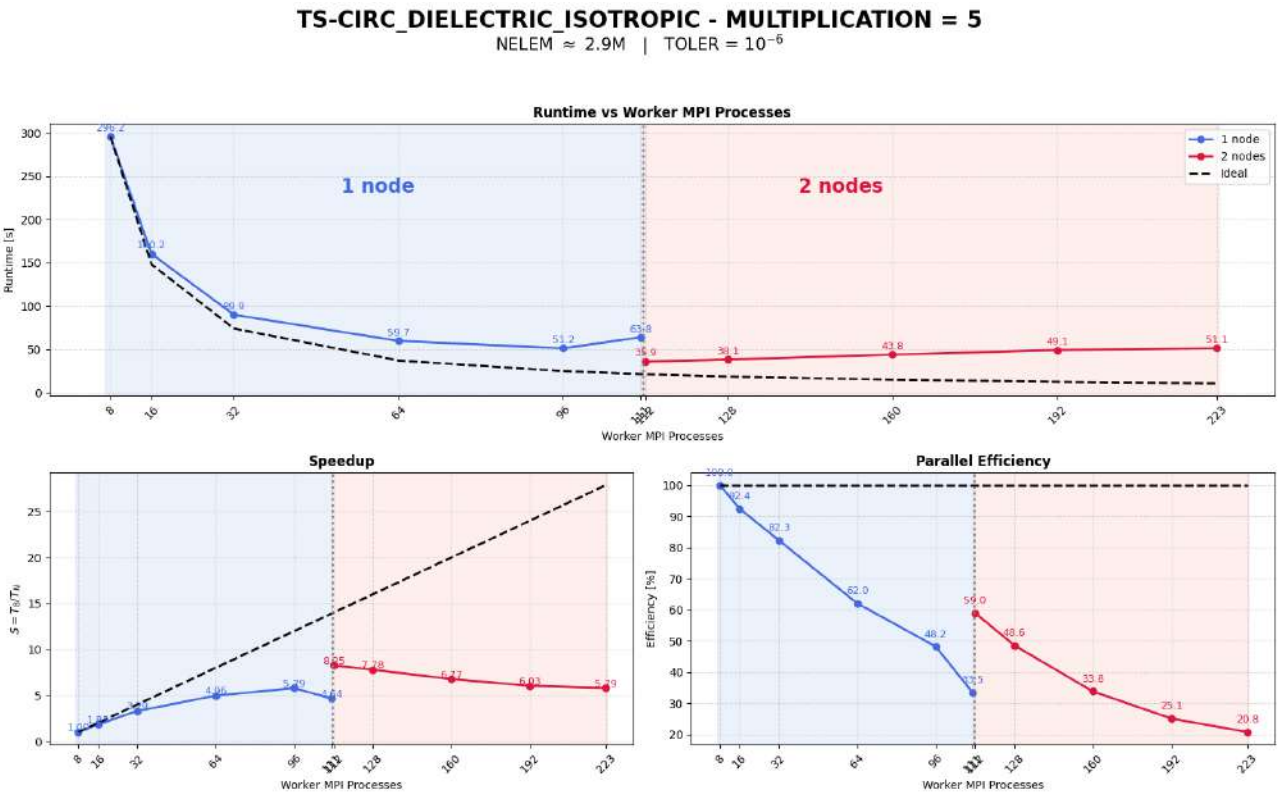

Figure 16: Strong-scaling results for the TS-CIRC_DIELECTRIC_ISOTROPIC case with approximately 2.9×10^6 elements and a solver tolerance of 10^{-6}. The blue and red regions correspond to one-node and two-node executions, respectively, while the dashed curves represent ideal strong scaling.

The first strong-scaling experiment was performed using the TS-CIRC_DIELECTRIC_ISOTROPIC test case. This is a two-dimensional electromagnetic problem discretized using first-order triangular finite elements. The computational domain consists of a circular dielectric region surrounded by vacuum, with an absorbing boundary condition imposed at the external boundary.

The original test case contains 2870 elements. For the scalability analysis, this mesh was progressively refined using the mesh multiplication functionality available in Alya. The multiplication factor controls the refinement of the original mesh, increasing the number of elements while preserving the geometry and physical configuration of the problem. A multiplication factor of 5 was used for the final scaling experiment, resulting in a mesh containing approximately 2.9×10^6 linear triangular elements and 1 471 457 degrees of freedom.

The use of a substantially finer mesh than in the original regression test was necessary for the strong-scaling analysis. If the numerical problem is too small, the computational work assigned to each MPI process becomes small even at relatively low core counts. Communication, synchronization and other parallel overheads can then dominate the execution time from the beginning of the scaling experiment. Increasing the mesh size therefore provides sufficient computational work per MPI process to obtain a more meaningful characterization of the parallel behaviour of EMWAVE.

The solver tolerance was fixed at 10^{-6} throughout the scaling experiment. Figure 16 shows the measured runtime, speedup and parallel efficiency as the number of CPU cores is increased. Since one MPI process was assigned to each core, the number of cores also corresponds to the number of parallel subdomains used in the simulation.

6.2.2 TS-CUBE_COLDPLASMA_ABC

The second strong-scaling experiment was performed using the `TS-CUBE_COLDPLASMA_ABC` test case. This configuration differs substantially from the previous benchmark. It represents a three-dimensional electromagnetic problem discretized using first-order tetrahedral Nédélec edge elements, whereas the `TS-CIRC_DIELECTRIC_ISOTROPIC` case is two-dimensional and uses linear triangular nodal elements. The material response is described using the anisotropic cold-plasma dielectric tensor, and absorbing boundary conditions (ABC) are imposed at the external boundaries of the computational domain.

The original mesh, corresponding to a multiplication factor of zero, contains 9,625 tetrahedral elements. For the same reason discussed for the previous benchmark, the mesh was progressively refined in order to increase the computational workload. In three dimensions, each multiplication level increases the number of tetrahedral elements by a factor of eight. The successive mesh sizes are therefore

$$9\,625 \;\rightarrow\; 77\,000 \;\rightarrow\; 616\,000 \;\rightarrow\; 4\,928\,000, \tag{6.4}$$

corresponding to multiplication factors from 0 to 3. The final strong-scaling experiment was performed with `MULTIPLICATION = 3`.

The physical parameters, material properties, boundary conditions and numerical formulation were kept unchanged throughout the scaling experiment. As in the previous scaling test, the solver tolerance was fixed at 10^{-6}, and the electromagnetic output was projected onto the original mesh. The corresponding strong-scaling results are presented in Figure 17.

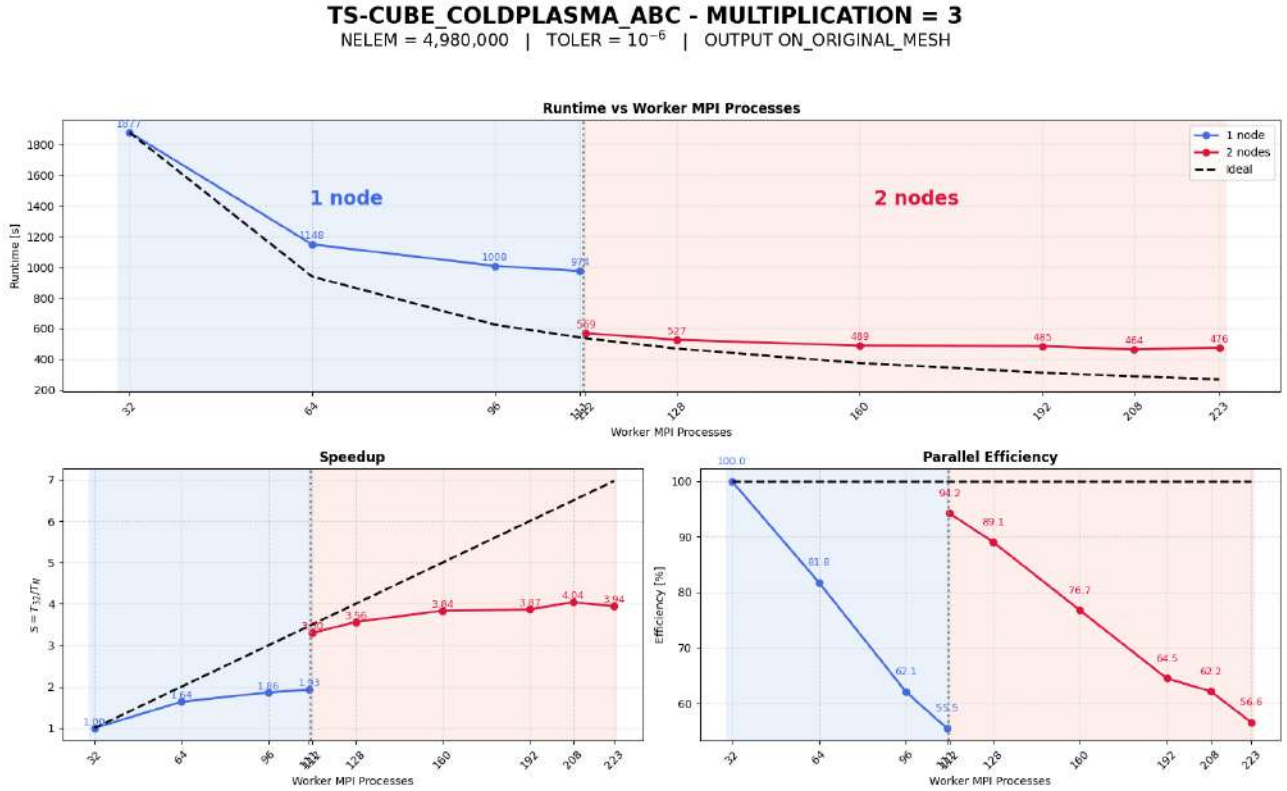

Figure 17: Strong-scaling results for the `TS-CUBE_COLDPLASMA_ABC` case with approximately 4.93×10^6 first-order tetrahedral elements and a solver tolerance of 10^{-6}. The blue and red regions correspond to one-node and two-node executions, respectively, while the dashed curves represent ideal strong scaling.

Despite the substantial differences between the two benchmark problems, the results reproduce the same main scaling characteristics observed for the `TS-CIRC_DIELECTRIC_ISOTROPIC` case. Within a single compute

node, the runtime decreases as the number of CPU cores is increased, although the corresponding speedup progressively departs from the ideal strong-scaling curve.

A pronounced discontinuity is again observed between 111 and 112 working cores, coinciding with the transition from a one-node to a two-node execution. The runtime decreases from approximately 974 s at 111 cores to 569 s at 112 cores, despite the addition of only one working core.

Beyond this transition, the speedup becomes almost flat. Increasing the number of working cores from 112 to 223, and therefore almost doubling the computational resources, reduces the runtime only from approximately 569 s to 476 s. The corresponding speedup increases only from approximately 3.30 to 3.94, while remaining far below the ideal strong-scaling curve.

The reproduction of these features is particularly relevant because the two benchmarks differ in dimensionality, geometry, mesh, material model and finite-element formulation. The first case is a two-dimensional isotropic dielectric problem solved using triangular nodal elements, whereas the present case is a three-dimensional cold-plasma problem solved using tetrahedral Nédélec edge elements. Nevertheless, both experiments exhibit a clear discontinuity at the one-node to two-node transition and a strong saturation of the speedup at larger core counts.

The agreement between the qualitative scaling behaviour of these two independent configurations therefore indicates that the observed limitations are not specific to a particular mesh or physical test case. Instead, they suggest a more general limitation of the current strong-scaling behaviour of the `EMWAVE` module.

6.2.3 Discussion

The strong-scaling experiments reveal two characteristic features of the current parallel behaviour of `EMWAVE`. First, the speedup progressively departs from the ideal strong-scaling curve and tends to saturate as the number of CPU cores is increased. Second, a pronounced discontinuity is observed between 111 and 112 working cores, corresponding to the transition from a one-node to a two-node execution.

The saturation of the speedup indicates that, for the problem sizes considered in this work, `EMWAVE` currently exhibits limited strong scalability. As the number of MPI processes increases while the problem size remains fixed, the computational workload assigned to each process becomes smaller. At the same time, communication, synchronization, data exchange at subdomain interfaces, solver operations and other parallel overheads become increasingly important relative to the local computational work. As a consequence, the use of additional CPU cores eventually provides diminishing returns. For the `TS-CUBE_COLDPLASMA_ABC` case, this results in only a small additional reduction in execution time at larger core counts. For the `TS-CIRC_DIELECTRIC_ISOTROPIC` case, the effect is even more pronounced: beyond 112 cores, the execution time increases as additional computational resources are employed. In this regime, the additional parallel overhead therefore outweighs the reduction in the local computational workload.

The discontinuity between 111 and 112 cores represents a different feature of the results. In both experiments, the transition to the second compute node produces a substantial reduction in execution time. The additional node changes several aspects of the execution simultaneously, including the distribution of the computational domain and the hardware resources available to the simulation. The present results do not allow to isolate the origin of this behaviour, and a dedicated profiling analysis would be required to determine whether the observed performance change originates mainly from the solver or from the parallel execution.

A particularly relevant result of the present analysis is the reproducibility of both characteristics in two substantially different electromagnetic problems. The `TS-CIRC_DIELECTRIC_ISOTROPIC` benchmark is a two-dimensional isotropic dielectric problem discretized using linear triangular nodal elements, whereas `TS-CUBE_COLDPLASMA_ABC` is a three-dimensional anisotropic cold-plasma problem discretized using first-order tetrahedral Nédélec edge elements. The two cases also employ different geometries, meshes and boundary conditions, with approximately 2.9×10^6 and 4.93×10^6 elements, respectively.

Despite these differences, both experiments reproduce the same qualitative scaling behaviour: a pronounced reduction in runtime at the transition from 111 to 112 cores and a strong saturation of the speedup at larger core counts. The fact that these features are recovered using different dimensionalities, element types and electromagnetic formulations suggests that they are not specific to a particular mesh or test configuration. Instead, they appear to characterize more generally the current parallel behaviour of the `EMWAVE` module.

These results should not be interpreted as a limitation of the parallel scalability of Alya as a whole. Alya has been developed for large-scale parallel simulations [19] and has been proven to have excellent scalability in up to 100 000 processors in certain cases [11], while `EMWAVE` is a comparatively recent module whose numerical and parallel implementation is still undergoing development and optimization. The present scaling campaign therefore provides a useful characterization of its current performance and identifies strong scaling as an area in

which further investigation and optimization could provide significant computational benefits.

A natural continuation of this work would consist of profiling the different stages of the `EMWAVE` solution process, including the linear solver, MPI communication and domain decomposition. Such an analysis would help identify the mechanisms responsible for both the saturation of the speedup and the reproducible performance change observed at the transition between one-node and two-node executions.

7 Sensitivity Analysis with DAKOTA

The electromagnetic response predicted by a numerical model depends not only on its mathematical formulation and spatial discretization, but also on the physical parameters used to define the problem. Assessing how the numerical solution changes when these parameters are varied provides useful information on the sensitivity and robustness of the simulated electromagnetic response.

In this work, a frequency sensitivity analysis was performed for the `EMWAVE` module using the DAKOTA [13] toolkit. DAKOTA was coupled to Alya in order to systematically modify the excitation frequency, execute the corresponding electromagnetic simulations, and collect the resulting electric-field profiles. The analysis was performed on the `TS-CUBE_DIELECTRIC_PEC_DAKOTA` test case, using a nominal frequency of 60 MHz.

Particular attention was given to the effect of the spatial discretization on the sensitivity results. An initial campaign performed on the original mesh revealed that the solution was insufficiently resolved to provide a reliable characterization of its response to frequency variations. The analysis was therefore repeated on a substantially refined mesh using MareNostrum 5. This comparison provides both a characterization of the frequency dependence of the solution and an illustration of the importance of adequate spatial resolution when performing numerical sensitivity studies.

7.1 Sensitivity Analysis and the DAKOTA Toolkit

Sensitivity analysis is used to investigate how variations in the input parameters of a numerical model affect its predicted response. Rather than considering only a single set of input values, selected parameters are systematically varied and the corresponding changes in one or more quantities of interest are examined. This provides information on the robustness of the numerical solution and on the degree to which the predicted response depends on the parameters used to define the problem.

In electromagnetic simulations, this type of analysis can be particularly useful because the solution may depend strongly on physical parameters such as the excitation frequency or the electromagnetic properties of the materials. A sensitivity study can therefore help distinguish parameters that have a limited influence on the solution from those for which relatively small variations produce significant changes in the computed electromagnetic fields.

The sensitivity analysis performed in this work was automated using DAKOTA (Design Analysis Kit for Optimization and Terascale Applications), a software toolkit developed at Sandia National Laboratories for the analysis of computational models [13]. DAKOTA provides methods for sensitivity and parameter studies, uncertainty quantification, parameter estimation, and numerical optimization. A particularly useful feature is its ability to interface with external simulation codes, allowing a numerical model to be repeatedly evaluated for different combinations of input parameters.

In the present work, DAKOTA was coupled to Alya/`EMWAVE` to investigate the sensitivity of the electromagnetic solution to variations of the excitation frequency. DAKOTA was therefore not used as an alternative electromagnetic solver: `EMWAVE` continued to solve the finite-element electromagnetic problem, while DAKOTA controlled the systematic variation of the selected input parameter and the execution of the corresponding simulations.

7.2 Coupling DAKOTA with Alya/EMWAVE

The coupling between DAKOTA and Alya/`EMWAVE` was implemented through an external-driver workflow. For each evaluation, DAKOTA generates a value of the parameter being investigated and provides it to the external driver. The driver then prepares the corresponding `EMWAVE` input, launches the electromagnetic simulation, and collects the quantity required for the subsequent analysis.

In the frequency sensitivity study considered here, the parameter controlled by DAKOTA is the excitation frequency. Starting from the reference `EMWAVE` configuration, the driver replaces the nominal frequency with the value associated with the current DAKOTA evaluation while leaving the remaining physical and numerical parameters unchanged. Alya is then executed for the modified configuration and `EMWAVE` computes the corresponding electromagnetic field.

This procedure is repeated automatically for all frequency samples defined in the DAKOTA study. The resulting electric-field solutions are then collected and compared along the same spatial sampling line, allowing the variation of both the field amplitude and its spatial profile to be evaluated as a function of the excitation frequency.

The overall coupling can therefore be summarized as

DAKOTA $\longrightarrow$ input parameter $\longrightarrow$ external driver $\longrightarrow$ EMWAVE simulation
$\longrightarrow$ selected response $\longrightarrow$ DAKOTA.

After each simulation, the external driver extracts the selected response quantity from the `EMWAVE` output and returns it to DAKOTA, which associates the result with the corresponding parameter evaluation. This approach has the advantage that the electromagnetic solver itself does not need to be modified for each individual parameter evaluation.

7.3 Reference Electromagnetic Problem

The sensitivity analysis was performed using the `TS-CUBE_DIELECTRIC_PEC_DAKOTA` test case. The original configuration is a three-dimensional problem discretized with 9625 linear tetrahedral elements and 1970 mesh nodes. The electromagnetic problem is solved using the edge-element formulation of `EMWAVE`, with the degrees of freedom associated with the edges of the tetrahedral mesh. This original discretization corresponds to `MULTIPLICATION`=0.

The computational domain consists of a cubic region containing a central isotropic dielectric, with relative permittivity $\varepsilon_r = 4$, embedded in vacuum. The relative permeability is $\mu_r = 1$ throughout the domain, and the materials are lossless. No plasma model is activated.

Perfect electric conductor (PEC) boundary conditions are imposed on the external boundary of the domain, constraining the tangential component of the electric field to zero. The electromagnetic field is excited by a prescribed current source oriented along the z direction, with a real amplitude of 50 A/m^2 and zero imaginary component.

The nominal excitation frequency is $f_0 = 60$ MHz. This configuration constitutes the reference problem for the sensitivity analysis: the excitation frequency is varied while the remaining parameters are kept unchanged.

7.4 Initial Sensitivity Analysis on the Original Mesh

The first sensitivity analysis was performed using the `TS-CUBE_DIELECTRIC_PEC_DAKOTA` case's original discretization, corresponding to `MULTIPLICATION`=0 and containing 9625 tetrahedral elements.

Starting from the nominal value

$$f_0 = 60 \text{ MHz}, \tag{7.1}$$

the frequency was varied by $\pm 20\%$, covering the interval

$$48 \text{ MHz} \leq f \leq 72 \text{ MHz}. \tag{7.2}$$

A total of 21 frequency values were considered. For each value, DAKOTA generated the corresponding parameter set and the coupled workflow automatically executed a new `EMWAVE` simulation. The resulting electric-field magnitude was then extracted along a diagonal crossing the entire domain, allowing the effect of the frequency variation on both the amplitude and spatial distribution of the solution to be examined.

The resulting electric-field profiles are shown in Figure 18. Despite the variation of the excitation frequency, the simulations retain a broadly similar spatial structure, with the electric-field magnitude increasing towards the centre of the domain and decreasing towards the external boundaries. The effect of the frequency variation is, however, strongly dependent on the spatial position.

Away from the centre, the different profiles remain relatively close to one another. In contrast, the largest dispersion is observed around the central field peak, as also indicated by the widening of the $\pm 1\sigma$ band. To quantify this dispersion, let

$$E_i(s) = |\mathbf{E}(s; f_i)| \tag{7.3}$$

denote the electric-field magnitude at position s along the diagonal for the i-th frequency sample. At each spatial position, the corresponding standard deviation over the $N = 21$ simulations is

$$\sigma_E(s) = \sqrt{\frac{1}{N} \sum_{i=1}^{N} \left[E_i(s) - \overline{E}(s) \right]^2}. \tag{7.4}$$

The $\pm 1\sigma$ band shown in Figure 18 therefore corresponds to the interval $\overline{E}(s) - \sigma_E(s)$ to $\overline{E}(s) + \sigma_E(s)$ at each position along the diagonal. Its width provides a local measure of the dispersion of the electromagnetic

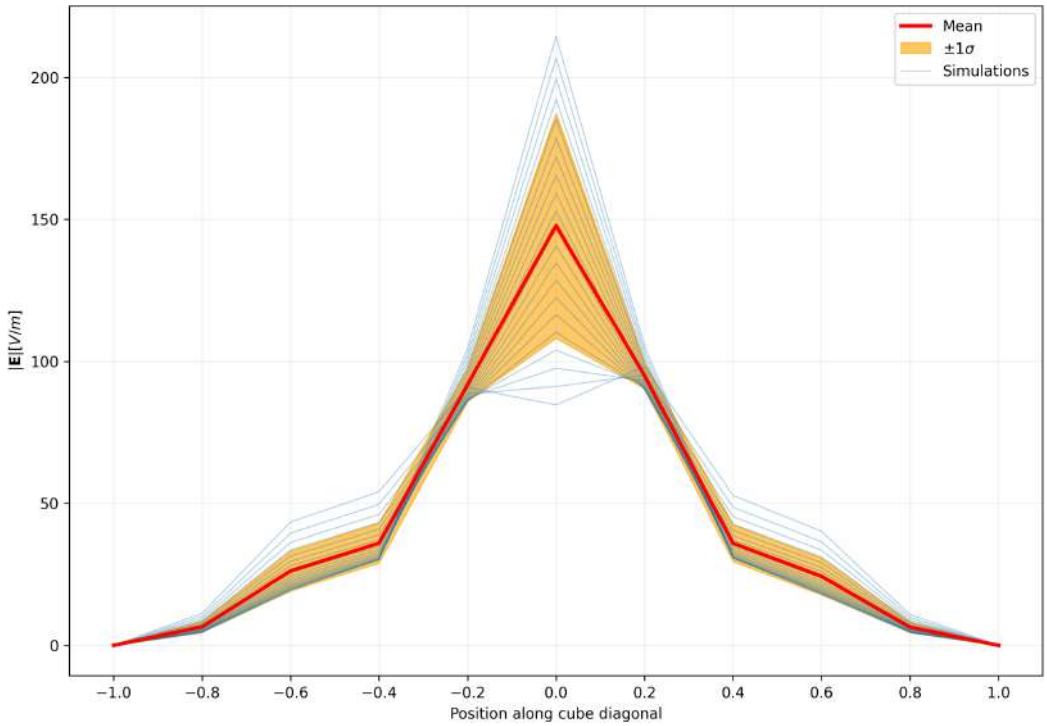

Figure 18: Electric-field magnitude along the cube diagonal for the 21 frequency samples of the initial sensitivity analysis using `MULTIPLICATION`=0. The individual simulations are shown together with their mean profile and the $\pm 1\sigma$ interval. The largest variability is concentrated around the central field peak.

response among the sampled frequencies. The larger width around the central peak therefore indicates that this region exhibits the strongest variation with excitation frequency.

The peak amplitude changes substantially across the frequency range, while some of the individual profiles exhibit a noticeably weaker central maximum. The sensitivity to frequency is therefore concentrated primarily in the region where the electric field reaches its highest amplitude.

At the nominal frequency of 60 MHz, the original mesh produces a central electric-field magnitude of approximately

$$|\mathbf{E}|_{\text{centre}} \simeq 146.8 \text{ V/m}. \tag{7.5}$$

Although this first campaign clearly showed a dependence of the electromagnetic response on the excitation frequency, the behaviour of the central peak also raised the question of whether the original mesh provided sufficient spatial resolution in the region where the solution varies most strongly. In particular, the relatively coarse discretization could make the predicted peak amplitude excessively sensitive to changes in the excitation frequency.

This observation motivated a second sensitivity campaign using a substantially refined mesh. The purpose of the refinement was therefore to determine whether the variability observed in Figure 18 represented the physical frequency sensitivity of the electromagnetic problem or was significantly affected by the spatial discretization.

7.5 Sensitivity Analysis on the Refined Mesh

To assess the influence of the spatial discretization on the sensitivity results, the frequency analysis was repeated using the same `TS-CUBE_DIELECTRIC_PEC_DAKOTA` configuration with `MULTIPLICATION`=2. The geometry, material properties, boundary conditions, and current source were left unchanged, so that the main difference with respect to the initial campaign was the spatial resolution of the electromagnetic problem.

The underlying input mesh remained the same, containing 1970 nodal points and 9625 tetrahedral elements. However, setting `MULTIPLICATION`=2 activates the mesh refinement procedure used by Alya, producing a substantially finer effective discretization of the same computational domain. The resulting increase in computational cost made it necessary to perform this second sensitivity campaign on MareNostrum 5.

The nominal excitation frequency was again

$$f_0 = 60 \text{ MHz}. \tag{7.6}$$

For the refined analysis, the perturbation interval was restricted to $\pm 10\%$ around the nominal value,

$$54 \text{ MHz} \leq f \leq 66 \text{ MHz}, \tag{7.7}$$

and a total of 21 frequency samples were considered. The resulting electric-field profiles along the cube diagonal are shown in Figure 19.

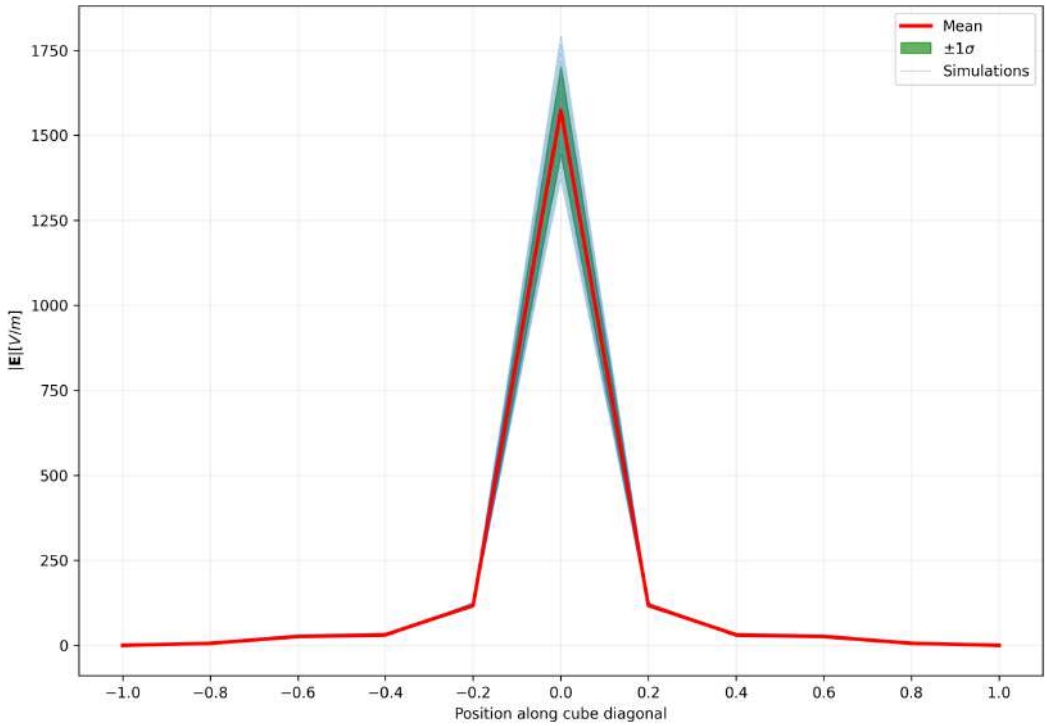

Figure 19: Electric-field magnitude along the cube diagonal for the 21 frequency samples of the refined sensitivity analysis using `MULTIPLICATION`=2. The individual simulations are shown together with their mean profile and the $\pm 1\sigma$ interval.

As in the initial analysis, the electric field reaches its largest amplitude near the centre of the cube and decreases towards the external boundaries. The refined calculations, however, reveal a substantially different central field amplitude.

The refined discretization gives

$$|\mathbf{E}|_{\text{centre}}^{\text{MULT=2}} \simeq 1568.4 \text{ V/m}. \tag{7.8}$$

The ratio between the two values is therefore

$$\frac{|\mathbf{E}|_{\text{centre}}^{\text{MULT=2}}}{|\mathbf{E}|_{\text{centre}}^{\text{MULT=0}}} \simeq 10.7. \tag{7.9}$$

The predicted central field amplitude therefore increases by more than one order of magnitude when the spatial discretization is refined. This large difference indicates that the original mesh was too coarse to adequately resolve the strong field peak occurring near the centre of the domain. Consequently, the sensitivity observed in the initial mesh were strongly influenced by the insufficient spacial resolution.

The refined campaign also allows the remaining frequency sensitivity to be quantified. At the central sampling point, the mean electric-field magnitude over the 21 simulations is

$$\overline{|\mathbf{E}|} = 1573.51 \text{ V/m}, \tag{7.10}$$

with a standard deviation of

$$\sigma_{|\mathbf{E}|} = 129.02 \text{ V/m}. \tag{7.11}$$

The corresponding coefficient of variation is

$$CV = \frac{\sigma_{|\mathbf{E}|}}{\overline{|\mathbf{E}|}} \times 100 \simeq 8.2\%. \tag{7.12}$$

Thus, within the investigated frequency interval, the standard deviation of the central field amplitude corresponds to approximately 8.2% of its mean value. Figure 19 also shows that the individual solutions remain grouped around a common spatial profile: the position and overall shape of the central peak are preserved, while its amplitude varies with the excitation frequency.

The comparison between the two sensitivity campaigns therefore highlights the strong influence of spatial resolution on the interpretation of parameter variations. With `MULTIPLICATION`=0, the central field peak was insufficiently resolved and the resulting frequency dependence was strongly affected by the coarse discretization. With `MULTIPLICATION`=2, the peak is resolved much more clearly and the 21 simulations exhibit a consistent spatial behaviour, with a coefficient of variation of approximately 8.2% at the centre of the domain.

These results show that a sensitivity analysis should not be interpreted solely in terms of variations of the physical input parameters. The numerical resolution of the underlying model must first be sufficiently fine to represent the relevant features of the solution. In the present case, increasing the mesh resolution revealed that part of the behaviour observed in the initial campaign was associated with the coarse discretization rather than with frequency variation alone. The refined analysis therefore provides a more representative characterization of the frequency sensitivity of the `EMWAVE` solution.

8 Conclusions

This work investigated and extended the capabilities of the EMWAVE module integrated into Alya's framework, a code based on the finite-element method with high-performance capabilities. In particular, EMWAVE is dedicated to simulating electromagnetic wave propagation in fusion-relevant plasma configurations.

Several developments were implemented during this work. First, the finite element discretization available in EMWAVE was extended to include quadratic triangular and quadrilateral elements, in addition to the existing first-order triangular formulation. Their compatibility with the electromagnetic formulation was assessed using dielectric and cold-plasma test cases. In particular, comparisons of the electric-field distributions and profiles showed that the different discretizations reproduce the same main spatial features, while differences in field amplitude remain due to the difference in the meshes and in the numbers of degrees of freedom considered.

The material description was extended through the implementation of electrical conductivity. This implementation allows simulation of materials with a complex permittivity, thus paving the way for simulations with lossy materials. In addition, the post-processing capabilities of EMWAVE were expanded through the implementation of the calculation of the magnetic field **H**, magnetic flux density **B**, and current density **J**. These quantities were implemented for the nodal and edge-element formulations of EMWAVE and incorporated into the standard Alya post-processing framework.

The electromagnetic implementation was independently validated against ERMES 20.0 using several conductive waveguide configurations. Comparisons of the electric field, magnetic field, magnetic flux density, and current density showed close agreement between the two codes. Across the investigated cases, the normalized root-mean-square differences remained below approximately 0.32%, providing quantitative validation of both the electromagnetic solution and the additional field quantities implemented in EMWAVE.

The computational behaviour of EMWAVE was also investigated through strong-scaling experiments on MareNostrum 5. Large refined electromagnetic problems were executed using an increasing number of MPI processes. The results showed an initial reduction in execution time as the number of computational resources was increased, followed by a progressive saturation of the speedup. For the TS-CIRC_DIELECTRIC_ISOTROPIC case, the execution time eventually increased at larger core counts. The results also show that the corresponding speedup and parallel efficiency depart substantially from ideal strong scaling. A discontinuity was observed when moving from executions contained within one compute node to executions distributed across multiple nodes. The present scaling tests do not provide sufficient information to determine the origin of this behaviour, which would require a more detailed analysis of the parallel execution.

Finally, an automated coupling between DAKOTA and Alya/EMWAVE was developed to investigate the sensitivity of the electromagnetic solution to variations in the excitation frequency. This analysis showed a clear electromagnetic variation, especially around the central field maximum, for the investigated frequency ranges. At the same time, the original and refined meshes demonstrated that the predicted field amplitude itself can depend strongly on spatial resolution. The sensitivity study therefore emphasized that an adequate discretization must be established before variations in the computed solution can be interpreted solely as a response to changes in physical input parameters.

Future work should address both the numerical limitations identified in the present study and the extension of the physical model. A systematic mesh-convergence study would first be required to establish the spatial resolution necessary for reliable electromagnetic solutions, particularly before performing further sensitivity analyses. The parallel behaviour of EMWAVE also requires a dedicated profiling analysis to understand the deterioration of the strong-scaling performance observed at larger core counts. From the physical point of view, a natural continuation of this work would be to move beyond the cold-plasma approximation through the implementation of a hot-plasma description. Such an extension would allow finite-temperature and kinetic effects neglected by the cold-plasma dielectric tensor to be included, providing a more realistic description of radio-frequency wave propagation and absorption in fusion plasmas.

References

[1] IPCC, *Climate Change 2022: Mitigation of Climate Change*, Contribution of Working Group III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change, Cambridge University Press, Cambridge, UK and New York, NY, USA, 2022.

[2] ITER Organization, *External Heating Systems*, ITER. Available at: https://www.iter.org/machine/supporting-systems/external-heating-systems

[3] F. Warmer, J. Alguacil, D. Biek, T. Bogaarts, G. Bongiovì, V. Bykov, J. P. Catalán, R. K. Duligal, I. Fernández-Berceruelo, S. Giambrone, C. Hume, M. Hrecinuc, R. Kembleton, J. Lion, T. Lyytinen, J. A. Noguerón Valiente, I. Palermo, V. Queral, D. Rapisarda, W. J. Rutten, L. Sanchis, X. Sarasola, K. Sedlak, A. Snicker, D. Sosa, and F. R. Urgorri, "Overview of European efforts and advances in stellarator power plant studies," *Fusion Engineering and Design*, vol. 202, Art. no. 114386, 2024.

[4] J. P. Freidberg, *Plasma Physics and Fusion Energy*, Cambridge University Press, Cambridge, 2007.

[5] H.-S. Bosch and G. M. Hale, "Improved formulas for fusion cross-sections and thermal reactivities," *Nuclear Fusion*, vol. 32, no. 4, pp. 611–631, 1992.

[6] L. Spitzer, *Physics of Fully Ionized Gases*, 2nd ed., Interscience Publishers, New York, 1962.

[7] J. Ongena, A. Messiaen, Y. O. Kazakov, R. Koch, R. Ragona, V. Bobkov, K. Crombé, F. Durodié, M. Goniche, A. Krivska, E. Lerche, F. Louche, A. Lyssoivan, M. Vervier, D. Van Eester, M. Van Schoor, T. Wauters, J. Wright, and S. Wukitch, "Recent advances in physics and technology of ion cyclotron resonance heating in view of future fusion reactors," *Plasma Physics and Controlled Fusion*, vol. 59, no. 5, 2017.

[8] J.-M. Jin, *The Finite Element Method in Electromagnetics*, 2nd ed., John Wiley & Sons, New York, 2002.

[9] H. Domingo Ramos, A. Soba Pascual, D. Gallart Escolà, and M. J. Mantsinen, "EMWAVE: Simulation of EM wave propagation in plasma using FEM," in *11th BSC Doctoral Symposium*, Barcelona Supercomputing Center, Barcelona, 2024.

[10] The Efficient Engineer, "Understanding the Finite Element Method," 2021.

[11] M. Vázquez, G. Houzeaux, S. Koric, A. Artigues, J. Aguado-Sierra, R. Arís, D. Mira, H. Calmet, F. Cucchietti, H. Owen, A. Taha, E. D. Burness, J. M. Cela, and M. Valero, "Alya: Multiphysics engineering simulation toward exascale," *Journal of Computational Science*, vol. 14, pp. 15–27, 2016.

[12] J.-C. Nédélec, "Mixed finite elements in $\mathbb{R}^3$," *Numerische Mathematik*, vol. 35, pp. 315–341, 1980.

[13] B. M. Adams, W. J. Bohnhoff, K. R. Dalbey, M. S. Ebeida, J. P. Eddy, M. S. Eldred, R. W. Hooper, P. D. Hough, K. T. Hu, J. D. Jakeman, M. Khalil, K. A. Maupin, J. A. Monschke, E. M. Ridgway, A. A. Rushdi, J. A. Stephens, L. P. Swiler, and D. M. Vigil, *Dakota, A Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis*, Sandia National Laboratories, 2024.

[14] J. D. Jackson, *Classical Electrodynamics*, 3rd ed., John Wiley & Sons, New York, 1999.

[15] T. H. Stix, *Waves in Plasmas*, American Institute of Physics, New York, 1992.

[16] D. G. Swanson, *Plasma Waves*, 2nd ed., Institute of Physics Publishing, Bristol and Philadelphia, 2003.

[17] M. Brambilla, "Numerical simulation of ion cyclotron waves in tokamak plasmas," *Plasma Physics and Controlled Fusion*, vol. 41, no. 1, pp. 1–34, 1999.

[18] M. Vázquez, G. Houzeaux, S. Koric, A. Artigues, J. Aguado-Sierra, R. Arís, D. Mira, H. Calmet, F. Cucchietti, H. Owen, A. Taha, E. D. Burness, J. M. Cela, and M. Valero, "Alya: Multiphysics engineering simulation toward exascale," *Journal of Computational Science*, vol. 14, pp. 15–27, 2016.

[19] C. Moulinec, G. Houzeaux, R. Borrell, A. Quintanas-Corominas, G. Oyarzun, J. Grasset, G. Giuntoli, and M. Vázquez, "A Massively Parallel Multi-Scale FE^2 Framework for Multi-Trillion Degrees of Freedom Simulations," in *Proceedings of the Platform for Advanced Scientific Computing Conference (PASC '23)*, pp. 1–11, 2023.

[20] R. Otín, "ERMES 20.0: Open-source finite element tool for computational electromagnetics in the frequency domain," *Computer Physics Communications*, vol. 310, Art. no. 109521, 2025.

[21] Barcelona Supercomputing Center, "MareNostrum 5: System Overview," *BSC Support Knowledge Center*, accessed August 2026.

[22] A. Grama, A. Gupta, G. Karypis, and V. Kumar, *Introduction to Parallel Computing*, 2nd ed., Addison-Wesley, 2003.

[23] CIMNE, *GiD: The Personal Pre and Post Processor*, International Center for Numerical Methods in Engineering (CIMNE), Barcelona, Spain.

[24] U. Ayachit, *The ParaView Guide: A Parallel Visualization Application*, Kitware, Inc., 2015. ISBN: 978-1-930934-30-6.

Acknowledgements

I would like to thank all three of my supervisors Hernán Domingo Ramos, Alejandro Soba and Ezequiel Goldberg, for their continuous help and guidance during these six months of internship at BSC. Their availability, and willingness to answer every question or doubt, together with the open, relaxed, and welcoming working environment they have made, has made it possible for me to dive into the numerical aspects of plasma physics, and absorb as much as I could during my internship at BSC. This is an experience that I will never forget.
I am also grateful for the opportunity to join the Fusion Group at BSC. For this, I would like to thank Mervi Mantsinen, the leader of the Fusion Group, for giving me the opportunity to carry out my internship within the group and to work on this project, as well as Hernán, Alejandro and Ezequiel for facilitating my integration into the group. I would also like to acknowledge the financial support provided by FuseNet and the Erasmus+ programme, which contributed to making this research stay at BSC possible.
I would also like to thank my colleagues from the Fusion Group, for making my internship much more enjoyable, and for the countless lunch breaks spent together, that I've turned into Spanish lectures for myself. I have reached a decent proficiency in Spanish, and I owe it to you. I hope that our future endeavours will bring us together again someday.
Last, but surely not least, I would like to thank my family, and especially my parents, for their unconditional and ever-present support. Without it, could never have begun, let alone finish my master's degree, and pursue my dreams of helping the cause of fusion energy. To say that I am grateful to them would be an understatement, and for this reason, I dedicate this work to them. I would like to finish by dedicating these words to my family:

Grazie per aver sempre creduto in me, anche nei momenti più difficili. Insieme, siamo capaci di tutto.
Vi voglio un bene dell'anima.