Accelerating simulations of electromagnetic waves in hot, magnetized fusion plasmas

Budé, R.H.S.

Award date:
2019

Eindhoven University of Technology (TU/e)

Master's thesis

Science and Technology of Nuclear Fusion – Applied Physics – Electrical Engineering

---

Accelerating Simulations of Electromagnetic Waves in Hot, Magnetized Fusion Plasmas

---

Author:

R. H. S. (Rick) Budé

Supervisors:

prof. dr. ir. A. B. (Bart) Smolders (TU/e)
prof. dr. R. J. E. (Roger) Jaspers (TU/e)
dr. ir. J. (Jan) van Dijk (TU/e)

External Advisor:

dr. D. (Dirk) van Eester (ERM/KMS - Laboratory for plasma physics)

August 23, 2019

Abstract

Large fusion devices like JET and ITER have various heating methods at their disposal that allow the machine to reach fusion relevant plasma temperatures. One of the heating methods employed in either machine is Ion Cyclotron Resonance Heating (ICRH). In this scheme the plasma is heated by injection of radio waves at a frequency matching the ion cyclotron resonance frequency of one of the species at a desired location in the plasma. Understanding where the injected power is absorbed, and by which species, is important for operating the fusion device, preparation for experiments and analysis of experimental data.

Traditionally numerical methods that accurately describe the behavior of the ICRH waves have been extremely CPU- and memory intensive. These so-called all-orders codes make no assumption on the size of the wavelength, compared to the Larmor radius in the plasma. However, the all-orders approach results in large, completely filled matrices, which need to be inverted in order to obtain the solution for the electromagnetic fields. One of the most prominent benchmarks in the field, the AORSA code, requires for example 10.000's of CPU hours to solve a 2D simulation for a tokamak. For stellarators with complex 3D magnetic topologies, using the same method would require unfeasible amounts of CPU hours and system memory.

The source of these complications is the integro-differential nature of the wave equation in hot plasmas. Efforts have been made in order to treat the problem as a purely differential problem, which yields sparse, banded matrices that are much faster to invert. The main approach is the usage of a truncated Finite Larmor Radius (FLR) expansion. It is based on the assumption that the wavelength of the waves in the plasma is much larger than the Larmor radius: $\rho_L k_\perp \ll 1$. The drawback of this method is that the damping of the Ion Bernstein Wave (IBW) is often underestimated, as this wave typically has wavelengths comparable to the Larmor radius: $\rho_L k_\perp \approx 1$.

This work examines both approaches and proposes a model that is capable of producing accurate predictions of the IBW damping, while using significantly less computational resources than an all-orders model. In the process, one cold plasma model is constructed, and three different hot plasma models. It is shown that the new hot plasma model is capable of producing electric fields with less than 1% error compared to the all-orders results. The main conclusion is that semi-localized approaches can produce accurate results, and that full-domain solutions like the all-orders approach are not always necessary. Future research should be focused on increasing the robustness of the model, and expansion to 2D and 3D simulations.

Preface

This master's thesis concludes the master's phase of my education at Eindhoven University of Technology (TU/e). A triple degree in Electrical Engineering (Electromagnetics specialization track), Applied Physics (Plasma specialization track) and Science and Technology of Nuclear Fusion is pursued. Although it was not easy to merge the three studies into a single study plan, the TU/e has always been supportive of this idea and I was assisted by countless people throughout the years.

One of the main questions has always been whether it would be possible to find a single master's project that matched all three studies. A perfect fit would be fusion plasma heating, or fusion plasma diagnosing, using electromagnetic waves. This question was finally answered in August 2018, when my supervisor for fusion, Roger Jaspers, contacted Dirk van Eester from the Royal Military Acadamy (ERM/KMS) in Brussels. Dirk van Eester replied, mentioning that he had a challenging project on modeling fusion plasma heating using the ICRH heating scheme. It became clear that the current models require 10 000s of hours of computational time, and that attempts to speed up the calculations have not been successful without affecting the accuracy of some of the wave phenomena in the plasma. This topic grasped my attention from the start, and the results from the ensuing year is documented in this thesis.

I would like to thank my supervisors from the TU/e: Bart Smolders, Roger Jaspers and Jan van Dijk. I believe that our meetings were always productive and helped me to define a clear path for my project. The discussions about my findings and results always led to new insights, and new steps to take. A special word of thanks goes to Dirk van Eester, who through his own expertise with ICRH modeling provided a lot of useful background information. I could build on the base that you built, which provided a solid starting point. I would not have come nearly as far without your assistance.

I would also like to thank my friends and family for standing with me throughout my project. Gratitude goes towards my parents, Frank and Sandra, for their continuous support throughout my project. Also, thanks go to my brother and sister, Roel and Lyanne, who are also studying at the TU/e, for allowing me to ramble about my project against you.

Finally, I would like to thank Myrthe, the girl that I love. You were always there when I needed you. Especially in the final weeks before the deadline, your presence, care and support helped me stay motivated and cheerful.

Thank you all, I could not have done this without you.

Table of Contents

1 Introduction

Burning massive quantities of fossil fuels for easy heat and power generation has served humanity well since the dawn of the industrial revolution. However, every passing day it becomes more evident that we cannot continue this path. Even though there is still plenty of coal, oil and natural gas left for the coming decades [1], there is an even more pressing reason to stop burning fossil fuels: global warming. In order to reduce the impact and severity of climate change, worldwide greenhouse gas (GHG) emissions need to be brought down significantly, as covered in the findings of the International Panel on Climate Change (IPCC) [2].

Therefore, carbon-free energy sources need to be developed and deployed on a larger scale, as fossil fuel burning accounts for more than two-thirds of the worldwide anthropogenic greenhouse gas emissions [3]. Currently, nuclear power and hydropower, and to a lesser extent solar power and wind power, are the main sources of carbon free electricity [4]. Other sources are being developed that may become of importance in the second half of this century, like Gen-IV fission and nuclear fusion [5], the latter being the focus of this work.

Nuclear fusion, the process of combining light elements into heavier elements, is the power source of our sun. In its core chains of nuclear reactions, mainly the so-called proton-proton chain reaction [6], merges the light hydrogen isotope ^{1}H into the heavier helium isotope ^{4}He. In this process some mass is lost, which is converted into energy through Einstein's famous $E = mc^2$. The energy gained in such a fusion reaction is multiple MeV [6]; a million times more than the few eV a typical chemical reaction, like burning a methane molecule, releases.

Replicating this process on earth and tailoring it for energy production purposes has been a longstanding technological challenge for decades. The Lawson criterion [7], quickly reveals under which conditions a fusion reaction produces enough energy to compensate for energy losses. If $n\tau_E T$, the triple product between the density n, the confinement time τ_E and the temperature T, exceeds a fuel- and temperature-dependent threshold, the fusion reaction can be self-sustaining. This product has been plotted in Figure 1 for three different fuels.

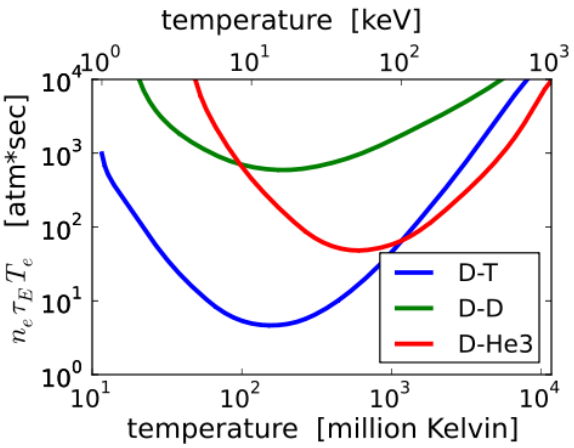

Figure 1: Triple product thresholds for three different fusion fuels: Deuterium+Tritium (D-T), Deuterium+Deuterium (D-D) and Deuterium+Helium-3 (D-He3). For a reaction to be energy-positive, $nT\tau$ must be higher than these fuel-dependent and temperature-dependent thresholds. (Image reprinted from Dstrozzi [8])

The lowest triple product threshold can be found for a D-T fuel mixture, formed by the heavy hydrogen isotopes deuterium (^{2}H or D) and tritium (^{3}H or T). Deuterium is abundant on earth, making up about 0.015% of seawater [9]. Tritium is a radioactive gas with a 12.3 years half-life [10] that can be produced inside a fusion reactor from lithium through breeding [11]. The fusion reaction for this fuel mixture is given in equation (1).

$$\mathrm{D} + \mathrm{T} \rightarrow {}^{4}\mathrm{He}\ (3.5\ \mathrm{MeV}) + \mathrm{n}\ (14.1\ \mathrm{MeV}). \quad (1)$$

Every fusion reaction releases a total of 17.6 MeV, carried as kinetic energy by the neutron and alpha (^{4}He) particle, which means that fusing 1kg of deuterium-tritium fuel mixture produces 189 GWh of energy. For reference, it would require about 23 000 tons of bituminous coal with a burn value of

30MJ/kg [12] to produce the same amount of energy. Additionally, in contrast to coal the fusion reaction releases no CO_2. Furthermore, there is no shortage of fuel supply in sight: both lithium and deuterium are available for the foreseeable future [13]. Finally, even though D-T fusion produces low- and intermediate-level radioactive waste through activation of structural materials by neutron irradiation and through tritium contamination, no long-lived high-level waste is produced [14].

Unfortunately, so far it has proven very difficult to build a fusion reactor that creates more power through fusion than it costs to run the machine. The ITER fusion reactor is a worldwide effort aimed at turning this around. It is currently under construction in Cadarache, France, designed to reach a thermal fusion output of 500MW with a thermal input of 50MW from auxiliary heating systems [15]. A scale model is shown in Figure 2.

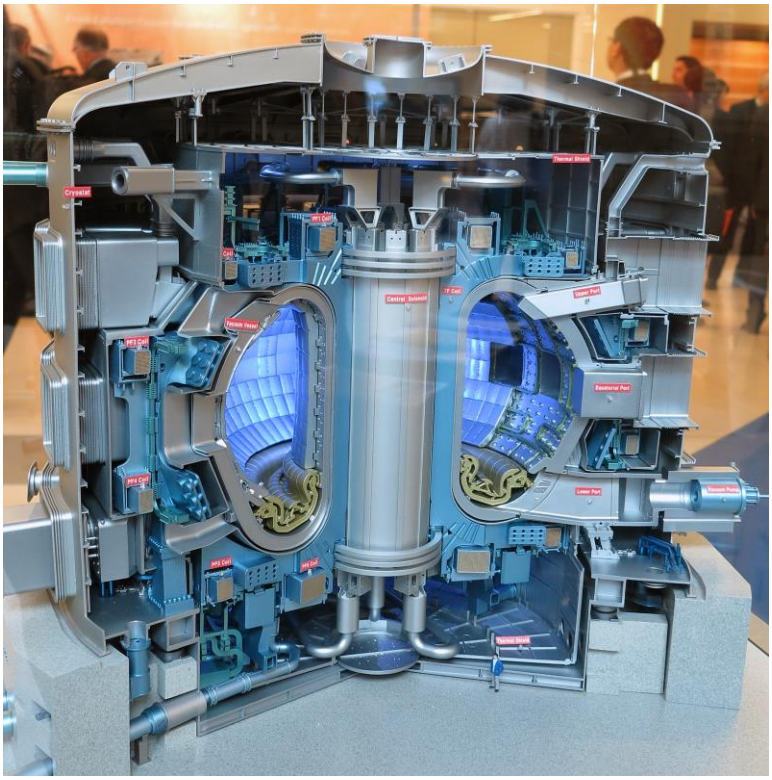

Figure 2: A scale model of the ITER reactor, exposed on the International Fusion Energy Days 2013. For scale, note the human figure standing on the ground. (Image adapted from IAEA imagebank [16])

The temperature of the fuel in ITER will be near the optimum temperature for D-T fusion, with particles having an energy of 14keV, or a temperature of about 160 million degrees (Figure 1). The fuel becomes a fully ionized plasma at these temperatures; all particles are charged and susceptible to magnetic and electric fields [17]. ITER therefore employs a tokamak design to obtain the complex magnetic topology required to contain the plasma. Tokamaks are a particular magnetic confinement fusion design, pioneered in the 1950s by Igor Tamm and Andrei Sakharov [18]. In this particular design, the fields are partly generated by external coils, and partly by driving a toroidal current in the plasma. Its relative simplicity made it a popular choice for fusion reactors. Some other examples of tokamaks are the Joint European Torus (JET), TEXTOR and ASDEX upgrade.

The confined plasma then needs to be brought up to ignition temperatures by auxiliary heating systems. Various heating schemes have been developed that can accomplish this, like Neutral Beam Injection (NBI), Electron Cyclotron Resonance Heating (ECRH) and Ion Cyclotron Resonance Heating (ICRH). The focus of this report is the ICRH scheme. Amongst others JET has ICRH at its disposal, and so will ITER. The advantage of this heating scheme is that the absorption location can easily be chosen by changing the antenna frequency. Furthermore, it directly heats ions in the plasma. Chapter 2.1 will discuss in more detail what ICRH is, and how it works.

However, there are some subtleties to wave propagation in hot fusion plasmas. Due to the presence of various absorption, reflection, cut-off and mode conversion layers in the plasma, it is difficult to determine where the injected power will end up exactly, and which species it will end up heating. Therefore, predicting where the heating power ends up nearly always requires numerical simulations. In addition, design of the ICRH antenna system requires deep insight in the physics of this heating strategy.

Finally, in preparation of experiments it is useful to design heating strategies upfront, to make sure the power ends up heating the desired species at the desired location.

Multiple numerical models have been created in the past for this purpose. Especially spectral codes like the All-Orders Spectral Algorithm (AORSA) and TOMCAT-U [19] produce accurate predictions of the RF electric field, and in many papers they are seen as a point-of-reference [20], [21], [22]. In these spectral codes the electric field is decomposed in a set of Fourier harmonics [23] or plane waves, as shown in equation (2) for a 2D simulation.

$$\vec{E}(x,y) = \sum_{n,m} \vec{E}_{n,m} e^{i(k_n x + k_m y)} \tag{2}$$

Unfortunately, a large computational cost is associated with this approach. With N, M and P modes in x, y and z direction respectively, a linear system of equations is obtained with $3 \times N \times M \times P$ equations and the same number of unknowns per equation [23]. A simple test is conducted that examines the time required to solve such a linear system, using randomly generated matrices of different sizes. The time required to invert such a dense matrix of complex doubles, as function of the number of rows N_{rows}, seems to scale approximately as $N_{rows}^{2.5}$, as demonstrated in Figure 3.

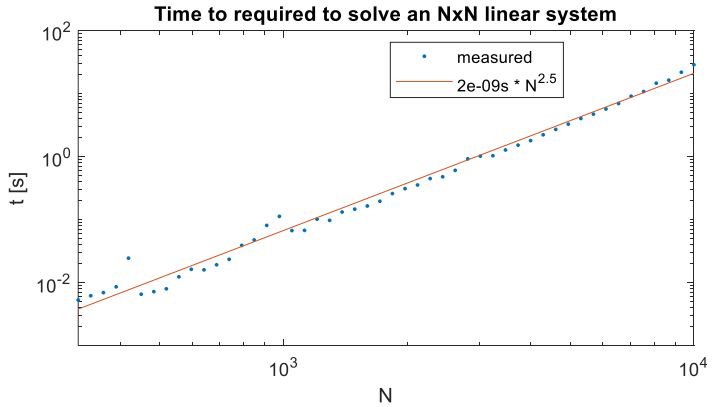

Figure 3: Time required to solve a linear system $Ax = b$, with A an $N_{rows} \times N_{rows}$ complex double matrix. The scaling seems to go approximately as $N^{2.5}$. Tests were executed in MATLAB on an i7-7700HQ laptop processor, limited to 4 computational threads.

Experimentation with the 1D AORSA-based all-orders model, discussed in chapter 5.1, suggests a minimum of 200 modes in each direction. Extrapolating the findings in Figure 3 for low-resolution 2D and 3D simulations yields Table 1. A 1D simulation will run just fine on any PC, but a 2D simulation is already stretching it for most workstations and/or laptops. A modern, powerful processor with a high core-count, and more than 250GB of RAM would be recommended. A 3D simulation, which would be desirable for stellarators with a more complex magnetic topology than tokamaks, can only be executed on a powerful supercomputer. At these timescales, experimental data, preparing for experiments, and performing parametric scans all become slow and computationally demanding tasks.

	N_{rows}	Storage space	Time
1D	600	5.8 MB	17 ms
2D	120 000	230 GB	2.8 hours
3D	24 000 000	9.2 PB	246 years

Table 1: Storage space requirements and time needed to solve the linear matrix system for a 1D, a 2D and a 3D scenario. The solve times have been extrapolated from Figure 3, under the assumptions that the same hardware is used and that the matrix fits in RAM.

The straightforward approach is adding more CPU cores to the problem, as AORSA parallelizes well [23]. Furthermore, some efforts have been made on running the AORSA matrix factorization on a GPU due to a GPU potentially having a better per-core performance and better power efficiency. Unfortunately, the initial result showed a data transfer bottleneck [24]. Other acceleration with hardware could come from high-speed storage, like Intel® Optane™ [25], which might benefit out-of-core solvers. It is even not unthinkable that hardware is designed, specifically created for inverting large matrices. In this work however, an attempt is made to reduce the amount of work required in the first place.

Some codes, like TORIC [21], EVE [26] and TOMCAT [27], partly tackles this problem by resorting to a truncated Finite Larmor Radius (FLR) approximation. In this approximation, it is assumed that the wavelength of interest is much larger than the Larmor radius, the radius with which particles gyrate in a magnetic field [19], throughout the domain. This results in sparse matrices that require significantly less storage space and computational power to invert [21]. For example, AORSA requires 8 hours on 576 cores to complete a 200x200 mode simulation [23], while TORIC can do it in just 9.5 CPU-hours [21].

Unfortunately, this approach breaks down for certain situations. Around the position where the ICRH waves are absorbed, Ion Bernstein Waves (IBW) can be generated through a process called mode conversion. The mode-converted IBW typically has wavelengths comparable to the Larmor radius. The assumption that the wavelength is much larger than the Larmor radius is violated by these mode-converted waves. As a result, the damping of these waves is not predicted in a satisfactory way. Figure 4 demonstrates this effect, using two of the hot plasma models that are discussed in chapter 5. The left figure is an all-orders model, based on AORSA. The right figure is a truncated Taylor series mode, bearing similarities with a truncated FLR code. Clearly, there are strong differences in the prediction of the short-wavelength IBW damping. Depending on the scenario, the differences can be quite large, leading to wrong conclusions about the amount of power absorption of the plasma, and the relative distribution of the absorbed power per species [19].

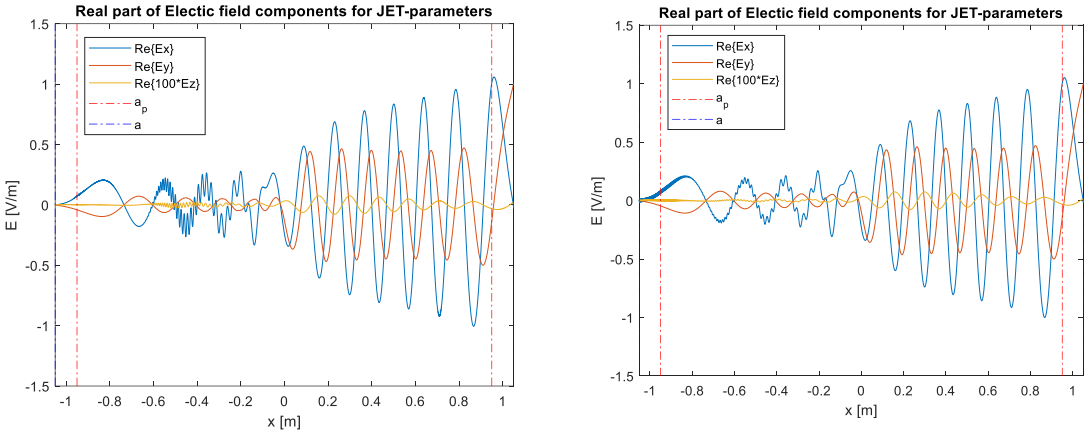

Figure 4: Simulation of the electric fields in JET-like plasma conditions. The left figure is created with the all-orders model, described in chapter 5.1, while the right figure is created with the truncated Taylor series model, described in chapter 5.2. For larger images, see Figure 19 and Figure 21. Note the amount of damping of the short-wavelength wave, to the left of $x = 0$.

This damping takes place exactly in the region where most of the power is absorbed. As such it is interesting to get accurate and fast calculations of what exactly happens in that region. This brings us to the research question for this work:

Can a model be developed that combines the accurate prediction for the Ion Bernstein Wave damping from the all-orders model with the computational speed of a truncated Finite Larmor Radius model?

A scheme is developed that comes close to this goal: it can approach the accuracy of an all-orders model, while having a computational complexity comparable to a truncated finite Larmor radius model.

Chapter 2 starts off with a discussion of wave behavior in cold and hot plasmas. Chapter 3 discusses some details related to the simulation setup. Chapter 4 introduces the cold plasma model, which is significantly easier to simulate than the hot plasma model. Chapter 5 introduces the three hot plasma models that were created, and they are compared with each other in Chapter 6. Chapter 7 discusses some peculiarities of the new scheme.

2 The physics of RF heating of hot, magnetized fusion plasmas

In a fusion reactor, the plasma can be described as a hot, magnetized fusion plasma. Section 2.1 and section 2.2 provide some information about the terms “hot” and “magnetized”. Section 2.3 continues with a description of the ICRH heating scheme, which is the central application for which all simulations are done. After this, a short review is provided about the time-harmonic Helmholtz equation, which is the core equation that the models solve. Finally, without going into too much depth, the dispersion relation for cold plasmas and hot plasmas are discussed, highlighting some of the features that will be visible in the simulations.

2.1 Differences between a hot and a cold plasma

In a thermodynamic equilibrium, all plasma particles have the same temperature. In weakly ionized, low density plasmas, this is not necessarily the case, with the electrons having a significantly higher temperature than the ions and the background gas: $T_e \gg T_i$. These plasmas are called non-thermal plasmas, or cold plasmas [28]. In a cold plasma, thermal motion is assumed to be negligible [29]. An example is a low-pressure DC glow discharge, for example in a fluorescent lamp, with typically $T_e \sim 10^4$K, while the ions and the background gas are approximately at room temperature [28].

In a hot plasma on the other hand, significant thermal motion can be observed. An example of a hot plasma near thermodynamic equilibrium is a fusion reactor plasma. The extremely high temperatures in the core lead to a full ionization of the plasma [17].

2.2 Effects specific to magnetized plasmas

When no background magnetic field is present, the plasma is an isotropic medium [30]. The properties are the same in all directions, which means that the waves do not have a preferred propagation direction. Anisotropy is introduced by adding a background magnetic field. When the magnetic field strength is high enough to significantly affect the movement of the particles in the plasma, one can speak of a magnetized plasma. A fusion plasma definitely falls in this category, but in nature it is also possible to find examples of magnetized plasmas, like astrophysical or ionospheric plasmas [30] [31]. In addition to anisotropy, the background magnetic field also adds extra inhomogeneity to the plasma.

Some phenomena are exclusive to magnetized plasmas. Thermal motion, for example, is forced into circular orbits, as discussed in section 2.3. Additionally, new groups of transverse waves become possible, called Alfvén waves [32]. Even for a cold plasma, the presence of a magnetic field already adds extra complexity, with more types of waves possible. In a hot plasma, this effect is enhanced due to the presence of the thermal motion of the ions.

2.3 The Ion Cyclotron Resonance Heating (ICRH) scheme

In a magnetic confinement fusion reactor, the plasma is confined with strong magnetic fields. A charged particle with a velocity $v_\perp$, perpendicular to this background magnetic field, will exhibit cyclotron motion due to the Lorentz force $\vec{F}_L$. Figure 5 provides a schematic drawing of this effect.

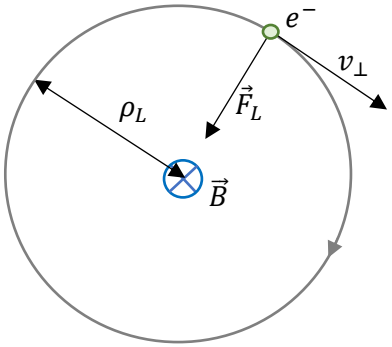

Figure 5: A schematic drawing of an electron exhibiting cyclotron motion. It "gyrates around a magnetic field line". The charged particle is kept in orbit by the Lorentz force.

This gyrating motion is called cyclotron motion. It has a characteristic orbit radius and a characteristic frequency. The radius of the circular motion is called the Larmor radius ρ_{Lj}. It is dependent on the velocity perpendicular to the magnetic field, for which often the thermal velocity v_j is assumed [23].

$$v_j = \sqrt{\frac{2k_B T}{m_j}} \tag{3}$$

$$\rho_{Lj} = \frac{v_j}{\Omega_{cj}} \tag{4}$$

The frequency with which the particles gyrate is the cyclotron resonance frequency Ω_{cj}. It depends on the mass m_j and the charge q_j of the species j of interest, and the background magnetic field strength $\left|\vec{B}(\vec{r})\right|$.

$$\Omega_{cj} = \frac{\left|q_j\right|\left|\vec{B}(\vec{r})\right|}{m_j} \tag{5}$$

The fundamental idea of ICRH is injecting RF waves into the fusion plasma with a frequency matching the ion cyclotron resonance frequency Ω_{cj} of the species j that should absorb the power. Important to notice is that $\left|\vec{B}(\vec{r})\right|$ is not uniform. Due to the setup of the machine, in general the background magnetic field is strongest near the core, and weaker near the edges of the machine following a $1/R$ decay. For the typical JET-like parameters as mentioned in Appendix B and the simplified magnetic field profile discussed in Chapter 3.2, the cyclotron frequency for hydrogen, Ω_{cH}, is plotted Figure 6. For reference, the antenna frequency ω is also indicated.

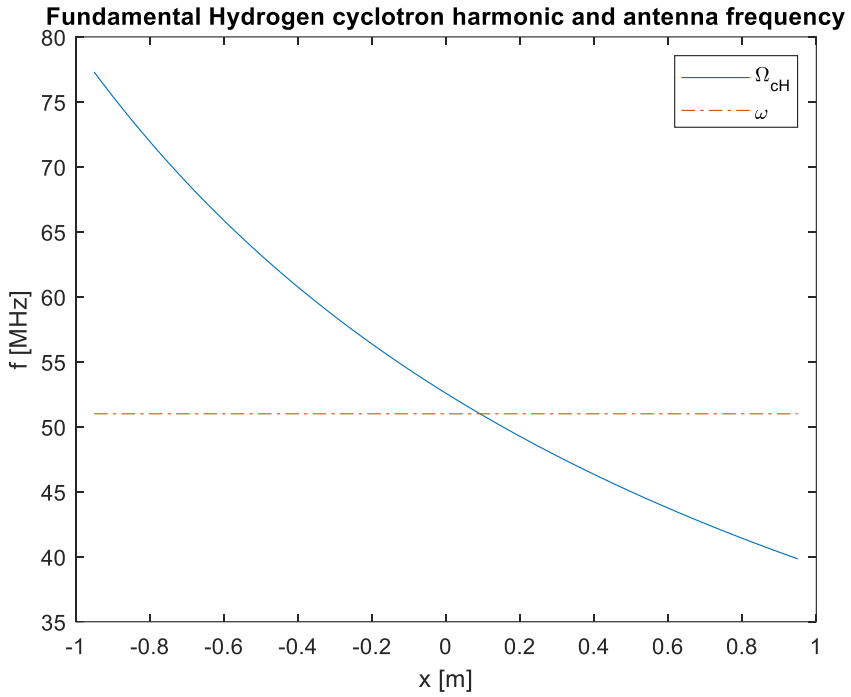

Figure 6: Graph of the fundamental hydrogen cyclotron frequency harmonic Ω_{cH} (in blue), and the antenna frequency ω (in orange) for typical JET parameters ($B_0 = 3.45T$, $f = 51MHz$). Absorption happens at the location where the antenna frequency matches one of the cyclotron harmonics of one of the species in the plasma.

Absorption primarily happens at the location where the antenna frequency matches one of the cyclotron harmonics of one of the species in the plasma. This is a very useful property of this heating scheme: the power deposition location can be selected by carefully tuning the antenna frequency.

In the case presented above, waves with a frequency of 51MHz are injected into the plasma. Near the core of the plasma, slightly to the right of $x = 0$, the antenna frequency matches the hydrogen ion cyclotron resonance frequency. Most power absorption will be centered around this location, as is visible in Figure 4. Worth noting is that in the JET case under consideration (Appendix B), only a 5% hydrogen concentration is present; the bulk of the plasma is deuterium. This demonstrates the effectiveness of the minority heating scheme, where minorities species like H or ^{3}He are injected into the plasma to establish efficient heating of the bulk plasma [33] [34].

2.4 The Helmholtz equation for time-harmonic fields in inhomogeneous media

In general, electromagnetic wave propagation is governed by the time-dependent Maxwell's equations. The differential form of these equations in free space is provided in (6)-(9). The electric field $\vec{E}$, the magnetic field $\vec{B}$, the current density $\vec{J}$ and the charge density ρ are all functions of the position $\vec{r}$ and time t [35].

$$\nabla \cdot \vec{E}(\vec{r}, t) = \frac{\rho(\vec{r}, t)}{\epsilon_0} \tag{6}$$

$$\nabla \cdot \vec{B}(\vec{r}, t) = 0 \tag{7}$$

$$\nabla \times \vec{E}(\vec{r}, t) = -\frac{\partial \vec{B}(\vec{r}, t)}{\partial t} \tag{8}$$

$$\nabla \times \vec{B}(\vec{r}, t) = \mu_0 \left(\vec{J}(\vec{r}, t) + \epsilon_0 \frac{\partial \vec{E}(\vec{r}, t)}{\partial t} \right) \tag{9}$$

When the Maxwell-Faraday equation (8) and Ampère's circuital law (9) are combined into a single equation, the time-dependent Helmholtz equation is obtained, eliminating $\vec{B}$ and ρ in the process.

$$\nabla \times \nabla \times \vec{E}(\vec{r}, t) + \epsilon_0 \mu_0 \frac{\partial^2 \vec{E}(\vec{r}, t)}{\partial t^2} = -\frac{\partial}{\partial t} \mu_0 \vec{J}(\vec{r}, t) \tag{10}$$

An extra simplification is made by assuming that both the electric field and the current is time-harmonic, which allows removal of the partial derivatives with respect to t. Additionally, the current density $\vec{J}$ is split into a contribution from the antenna $\vec{J}_{ant}$ and the response from the plasma $\vec{J}_p$, such that $\vec{J} = \vec{J}_{ant} + \vec{J}_p$. In inhomogeneous media, the time-harmonic Helmholtz equation is then given by: [23]

$$-\nabla \times \nabla \times \vec{E}(\vec{r}) + k_0^2 \vec{E}(\vec{r}) = -i\omega\mu_0 \left(\vec{J}_{ant}(\vec{r}) + \vec{J}_p(\vec{r}) \right). \tag{11}$$

Here, $k_0 = \omega/c$ is the free-space wavenumber. The antenna current density $\vec{J}_{ant}$ can be used for modeling the excitation of the plasma due to the ICRH antennas in the edge. The simplest possible model is that of a virtual, infinitely thin oscillating current sheet. However, this results in strong discontinuities in the derivatives of the electric field around this current sheath.

Much more well-behaved solutions are obtained by setting a finite value for one of the electric field components at a predefined position near the plasma edge. This effectively models the common situation where a Faraday screen is present in front of the ICRH antenna, preventing the electric field component parallel to the background magnetic field from passing through [36]. This eliminates $\vec{J}_{ant}$ from (11), yielding:

$$-\nabla \times \nabla \times \vec{E}(\vec{r}) + k_0^2 \vec{E}(\vec{r}) = -i\omega\mu_0 \vec{J}_p(\vec{r}). \tag{12}$$

The electric field $\vec{E}(\vec{r})$ is the desired output of the model. However, it requires a proper description of the plasma current $\vec{J}_p$. For cold plasmas it will turn out to be easy to describe the plasma current, but it will be significantly more difficult for the hot plasma case.

2.5 Electromagnetic waves in cold, magnetized plasmas

As a starting point for more advanced hot-plasma simulations, simulations are performed in a cold, magnetized plasma. Compared to the hot plasma, it is easier to formulate an expression for the plasma current density. This is for example demonstrated in textbooks from Stix [37] and Swanson [29]. Even though the hot plasma theory is more difficult, the cold plasma demonstrates how the plasma physics can be abstracted away in a dielectric tensor. Many components that are written for the cold plasma model can be reused for the hot plasma models. In fact, in all hot plasma models that are created, the hot plasma dielectric tensor can be swapped for the cold plasma dielectric tensor by changing a single line of code.

2.5.1 The cold plasma dielectric tensor

To show how behavior of particles in a plasma can be captured in the plasma current density through a dielectric tensor, an approach similar to D.G. Swanson's *Plasma Waves* is followed [29]. The analyses starts by noting that the equation of motion of a particle of species j, exhibiting time-harmonic oscillations in a cold, collisionless plasma, is written as:

$$-i\omega m_j \vec{v}_{1j} = q_j\left(\vec{E}_1 + \vec{v}_{1j} \times \vec{B}_0\right). \tag{13}$$

In this equation, the amplitude of the oscillation of the particles is given by $\vec{v}_{1j}$, the RF electric field amplitude is indicated with $\vec{E}_1$, and the static background magnetic field is given by $\vec{B}_0$. For an arbitrarily directed $\vec{B}_0$, equation (13) can be rewritten as a system of equations, linear in $\vec{v}_{1j}$.

$$\begin{bmatrix} -i\omega m_j & -B_{0,z}q_j & B_{0,y}q_j \\ B_{0,z}q_j & -i\omega m_j & -B_{0,x}q_j \\ -B_{0,y}q_j & B_{0,x} & -i\omega m_j \end{bmatrix} \begin{bmatrix} v_{1jx} \\ v_{1jy} \\ v_{1jz} \end{bmatrix} = \begin{bmatrix} q_j E_{1x} \\ q_j E_{1y} \\ q_j E_{1z} \end{bmatrix} \tag{14}$$

Often the background magnetic field $\vec{B}_0$ is assumed to be in the $\hat{z}$-direction, such that $\vec{B} = B_z\hat{z}$. In order to remain compatible with theory, this choice is made here as well. In Appendix A the derivation for arbitrarily directed $\vec{B}_0$ is provided. Solving the system for $\vec{v}_{j1}$ allows rewriting (14) as a system of equations, linear in $\vec{E}_1$. The definition for the ion cyclotron frequency (5) is applied for a more compact notation.

$$\begin{bmatrix} v_{1jx} \\ v_{1jy} \\ v_{1jz} \end{bmatrix} = \frac{1}{m_j^2\omega\left(\Omega_{cj}^2 - \omega^2\right)} \begin{bmatrix} -im_j\omega^2 q_j & B_{0,z}\omega q_j^2 & 0 \\ -B_{0,z}\omega q_j^2 & -im_j\omega^2 q_j & 0 \\ 0 & 0 & iq_j m_j\left(\Omega_{cj}^2 - \omega^2\right) \end{bmatrix} \begin{bmatrix} E_{1x} \\ E_{1y} \\ E_{1z} \end{bmatrix} \tag{15}$$

In general, the plasma current in a cold plasma is given in terms of the particle velocity $\vec{v}_j$. However, through the process described above it is demonstrated that the plasma current might as well be expressed in terms of the electric field $\vec{E}_1$. This allows capturing the plasma behavior in an abstract conductivity tensor $\vec{\vec{\sigma}}$.

$$\vec{J}_p = \sum_j n_j q_j \vec{v}_j = \vec{\vec{\sigma}}\vec{E} \tag{16}$$

$$\vec{\vec{\sigma}} = \sum_j \frac{n_j q_j}{m_j^2\omega\left(\Omega_{cj}^2 - \omega^2\right)} \begin{bmatrix} -im_j\omega^2 q_j & B_{0,z}\omega q_j^2 & 0 \\ -B_{0,z}\omega q_j^2 & -im_j\omega^2 q_j & 0 \\ 0 & 0 & iq_j m_j\left(\Omega_{cj}^2 - \omega^2\right) \end{bmatrix} \tag{17}$$

At this point the dielectric tensor is introduced as a means of simplifying the Helmholtz equation a bit further. In this expression, $\vec{\vec{I}}$ is a 3×3 unit dyadic.

$$\vec{\vec{\epsilon}}(x) = \vec{\vec{I}} + \frac{i}{\omega \varepsilon_0} \vec{\vec{\sigma}}(x) \tag{18}$$

Some extra steps can be taken to further simplify (15) and (17), leading to expressions for the dielectric tensor. However, these derivations can easily be found in literature, and involve only small mathematical manipulations. Therefore, only the final expression for the cold-plasma dielectric tensor is provided, assuming that $\vec{B}_0$ is pointed in the $\hat{z}$-direction [29].

$$\vec{\vec{\epsilon}} = \begin{pmatrix} S & -iD & 0 \\ iD & S & 0 \\ 0 & 0 & P \end{pmatrix} \tag{19}$$

S, D, R, L and P are the Stix cold plasma dielectric tensor terms. A collision frequency ν is inserted using a Krook model for both the ions and the electrons [38]. This adds a bit of damping to the model, and it makes dealing with the singularities at $\omega = \pm\Omega_{cj}$ numerically easier.

$$R = 1 - \sum_j \frac{\omega_{pj}^2}{\omega\left[\omega + i\nu_j + q_j \Omega_{cj}\right]} \tag{20}$$

$$L = 1 - \sum_j \frac{\omega_{pj}^2}{\omega\left[\omega + i\nu_j - q_j \Omega_{cj}\right]} \tag{21}$$

$$P = 1 - \sum_j \frac{\omega_{pj}^2}{\omega\left(\omega + i\nu_j\right)} \tag{22}$$

$$S = \frac{1}{2}(R + L) \tag{23}$$

$$D = \frac{1}{2}(R - L) \tag{24}$$

In these expressions, ω_{pj} denotes the plasma frequency, and q_j denotes the sign of the charge of particle species j. They are defined as:

$$\omega_{pj} = \sqrt{\frac{n_j q_j^2}{\epsilon_0 m_j}} \tag{25}$$

$$q_j = \frac{q_j}{|q_j|} \tag{26}$$

2.5.2 The cold plasma dispersion relation

In the previous section a definition for the dielectric tensor was provided. When this is substituted in the Helmholtz equation, equation (12), a very short expression remains for the cold plasma.

$$-\nabla \times \nabla \times \vec{E}(x) + k_0^2 \vec{\vec{\epsilon}}(x) \cdot \vec{E}(x) = 0 \tag{27}$$

The dispersion relation can be derived from here, describing the wavelengths that the plasma allows, in relation to the excitation frequency. It can be calculated for every point in the plasma, as if it were an infinitely extended homogeneous plasma, with a plane wave with wavevector $\vec{k}$ traveling through it. Under the plane wave approximation, the curl-curl can be substituted by outer products with the normalized wave vector $\vec{n}$: $\vec{n} = \vec{k}/k_0$. Solutions are now found for the amplitude and phase of the spectral electric field vector $\vec{E}_k$.

$$\vec{n} \times \vec{n} \times \vec{E}_k + \vec{\vec{\epsilon}} \cdot \vec{E}_k = 0 \tag{28}$$

The wavevector components parallel and perpendicular to the magnetic field are denoted as $n_\parallel$ and $n_\perp$ respectively. Without loss of generality, $\vec{B}$ can be fixed in the $\hat{z}$-direction, and $\vec{n}$ in the $x - z$ plane, as is proposed in Swanson [29]. Under these assumptions, $n_x = n_\perp$, $n_y = 0$ and $n_z = n_\parallel$. Using the expression for the cold plasma dielectric tensor (19), equation (28) can be written as:

$$\begin{vmatrix} S - n_\parallel^2 & -iD & n_\perp n_\parallel \\ iD & S - \left(n_\parallel^2 + n_\perp^2\right) & 0 \\ n_\perp n_\parallel & 0 & P - n_\perp^2 \end{vmatrix} \begin{bmatrix} E_x \\ E_y \\ E_z \end{bmatrix} = 0. \tag{29}$$

Finding the roots of this system delivers the dispersion relation. This process can be solved by setting the determinant of (29) to 0. For a cold plasma this can be done analytically. When a value is assumed for $k_\parallel$, for example based on experimental observations of the antenna spectrum, a fourth order polynomial of $n_\perp$ is obtained.

$$An_\perp^4 + Bn_\perp^2 + C = 0 \tag{30}$$

$$A = S \tag{31}$$

$$B = D^2 - S^2 - PS + (S + P)n_\parallel^2 \tag{32}$$

$$C = P\left(S^2 - D^2 - 2Sn_\parallel^2 + n_\parallel^4\right) \tag{33}$$

Finding the roots of (30) yields the perpendicular wavelength of waves that are allowed by the homogeneous plasma, as function of the magnetic background field strength $\vec{B}_0$, the particle densities n_j, the excitation frequency ω and the parallel wavenumber $k_\parallel$. In general S, P and D are complex due to the presence of the collision frequency ν. Therefore, finding the roots (30) yields four complex solutions for $k_\perp$, where the real parts indicate the wavelength, and the imaginary parts indicate damping.

The dispersion roots for the JET case under consideration are plotted in Figure 7. A modified logplot is used, plotting k_x^2 for $-1 < k_x^2 < 1$, and $\text{sign}(k_x^2) * \log_{10}(|k_x^2|)$ otherwise. The results have been grouped by the dispersion roots belonging to the fast wave and the slow wave. The fast wave has in general the largest wavelength of the two.

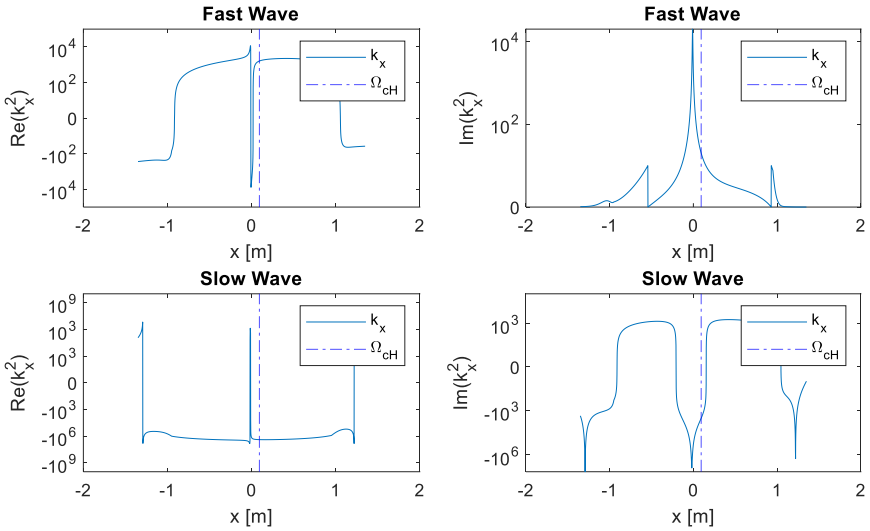

Figure 7: Modified logplot showing the solutions for the cold plasma dielectric tensor. The fast wave and slow wave dispersion roots for the JET case under consideration (Appendix B) are separated from each other. Throughout most of the simulation domain, the fast wave is propagative ($Re(k_x^2) > 0$), while the slow wave is strongly damped. Only near the plasma edge, and near the ion cyclotron resonance layer, the situation is different.

The fundamental hydrogen ion cyclotron resonance layer is found slightly to the right of the plasma center. Close to this location, the real part of the fast wave root (top left) dives down to negative values. In addition, there is a spike in the imaginary part (top right). Combined, this indicates a strongly damped fast wave. The slow wave is not propagative throughout most of the simulation domain; the real part of the slow wave root (bottom left) is nearly always negative.

For comparison a cold plasma simulation of the electric field using the same parameters is provided in Figure 8. The implementation details for this model are discussed in Chapter 4. The waves travel from right to left. Around $x = 0$, a strong decrease in the fast wave amplitude is visible, as expected from the dispersion plots. The slow wave makes an appearance in the non-physical spill-over areas at edges of the simulation domain, beyond $x = \pm a$, due to the low density at this location. As discussed in chapter 3.5, the spill-over regions are usually not considered in the simulations

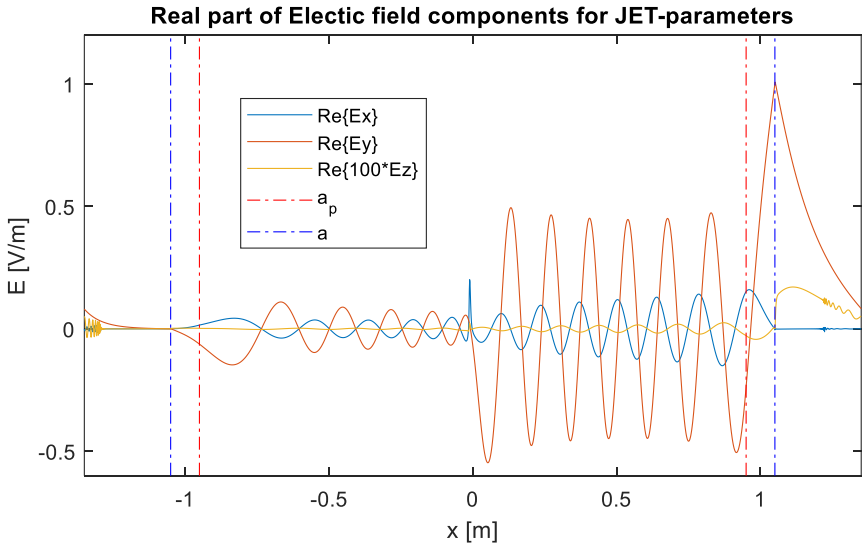

Figure 8: Cold-plasma results showing some of the wave characteristics. Notice the damping of the fast wave at the fundamental hydrogen ion cyclotron resonance layer near $x = 0$. Additionally, evanescent behavior of the fast wave is visible beyond $x = \pm a_p$. The slow wave makes an appearance in the non-physical spill-over regions at the edges of the simulation domain, beyond $x = \pm a$.

2.6 Electromagnetic waves in hot, magnetized plasmas

The wave description becomes more complicated for hot plasmas. The main difference is the fact that the particles now exhibit significant thermal motion. As mentioned in section 2.1, due to the background magnetic field the charged particles will exhibit cyclotron motion: they "gyrate around magnetic field lines". Aside from that, the particles have a temperature distribution. In this chapter, the dielectric tensor for a hot plasma is discussed, and some hot plasma effects are highlighted that differ from a cold plasma case.

2.6.1 The hot plasma dielectric tensor

The dielectric tensor for the cold plasma is a useful tool that allows abstracting away the plasma properties as a spatially depending dielectric tensor. This allows the usage of the regular Maxwell's equations for solving the wave problem. A similar result is desired for the hot plasma, as it allows reuse of the existing tooling and experience gained with the cold plasma model. The following section highlights some of the key steps in the derivation of the hot plasma dielectric tensor in chapter 4.3 from D.G. Swanson's *Plasma Waves* [39], concluding with its definition of the dielectric tensor for Maxwellian velocity distributions. Very similar, and older, derivations can be found in work from Stix [40], although this work unfortunately is not written in SI units.

Distributions of particles, both in the spatial domain and the velocity domain, become important for hot plasmas. This is for example apparent in the description of the plasma current, which in the hot plasma case requires integrals over the particle distribution functions in velocity space [41].

$$\vec{J}_p(\vec{r},t) = \sum_j q_j \int_v \vec{v} f_j(\vec{r},\vec{v},t) d^3v \tag{34}$$

This dependence on the distribution functions must be considered in the derivation of the dielectric tensor. In most treatments of the hot plasma dielectric tensor, the collisionless Boltzmann equation, in some works referred to as the Vlasov equation or the kinetic equation, is taken as the starting point [42]. It describes the evolution of f_j in the absence of collisions. Furthermore, the assumption is made here that only the Lorentz force is acting on the particles, which yields equation (35). q_j, m_j and f_j are the charge, mass and distribution function of species j respectively. When comparing this with (13), for the cold plasma dielectric tensor, the increased mathematical complexity in case of a hot plasma can be seen:

$$\frac{\partial f_j(\vec{r},\vec{v},t)}{\partial t} + \vec{v}\nabla f_j(\vec{r},\vec{v},t) + \frac{q_j}{m_j}\left[\vec{E}(\vec{r},t) + \vec{v}\times\vec{B}(\vec{r},t)\right]\nabla_v f_j(\vec{r},\vec{v},t) = 0. \tag{35}$$

Note how through the plasma current (34), the time-dependent Helmholtz equation (10) is strongly tied to the Vlasov equation (35). This is a strong motivation for finding a description of the dielectric tensor; as it can be used to abstract away this coupling.

Fortunately, in the RF-heated plasmas under consideration, there are two processes at different timescales. On the one hand, there is the RF time-harmonic perturbation of the plasma. On the other hand, there is a slow evolution of the distribution functions. Absorption of the waves by the plasma for example creates an energetic tail of particles that influences the wave propagation. These radically different timescales motivate a linearization of the Vlasov equation, as mentioned for example in Stix [40], Bittencourt [42] and van Eester and Lerche [43].

$$f_j(\vec{r},\vec{v},t) = f_{0j}(\vec{r},\vec{v},t) + f_{RFj}(\vec{r},\vec{v},t) \qquad \left(|f_{RFj}| \ll f_{0j}\right) \tag{36}$$

Here, f_{0j} contains the slowly varying, quasi-static distribution function, while f_{RFj} contains the RF perturbation.

In Swanson [39] it is than shown that through some manipulations, the RF distribution function f_{RFj} can be expressed in terms of the quasi-static f_{0j} by integrating along the zero-order particle trajectories in space, from time $t' = -\infty$ to $t' = t$ [40] [44]:

$$f_{RFj}(\vec{r},\vec{v},t) = -\frac{q_j}{m_j}\int_{-\infty}^{t}\left[\vec{E}_{RF}(\vec{r}',t') + \vec{v}' \times \vec{B}_{RF}(\vec{r}',t')\right]\cdot\nabla_{v'}f_{0j}(\vec{v}')dt'. \tag{37}$$

This is the first step to writing the RF response of the plasma in terms of a dielectric tensor. By making suitable approximations for the fields and the motion of the particles, the integral can be performed analytically. For example, for the velocity the circular gyrotron motion is assumed. Additionally, $\vec{E}_{RF}$ and $\vec{B}_{RF}$ are assumed to be time-harmonic.

The separation of timescales allows for an iterative scheme, where the wave equation is solved separately from the kinetic equation. The wave equation updates the RF electric field $\vec{E}_{RF}$, and the kinetic equation is used for updating the equilibrium distribution f_{0j} [43]. In this work the iteration is out of scope; only the wave equation is solved, as highlighted in Figure 9.

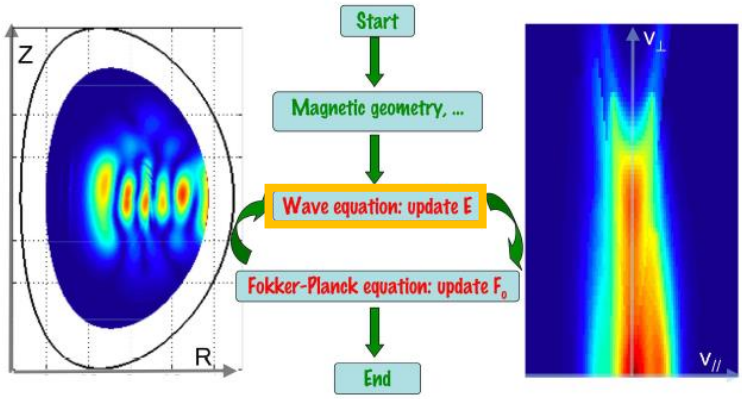

Figure 9: Iterative scheme enabled by linearizing the particle distribution contributions (reprinted and modified from [43]). The current project solely focuses on solving the wave equation for a single iteration, indicated with the orange box.

Combining (34) and (37) yields, through a lot of manipulations, an expression for the dielectric tensor that only depends on static, or semi-static plasma parameters. These include for example the slowly varying equilibrium distribution f_{0j}, the background magnetic field, the density and the excitation frequency. It requires integrals over the parallel and the perpendicular velocity dimensions, $v_\parallel$ and $v_\perp$ respectively [39]. As an example, ϵ_{33} is provided, the other eight components of $\vec{\vec{\epsilon}}$ have similar forms.

$$\epsilon_{33} = 1 + \sum_j \frac{\omega_{pj}}{\omega}\sum_{n=-\infty}^{\infty} 2\pi\int_{\infty}^{\infty} dv_\parallel \int_0^{\infty} dv_\perp \frac{nJ_n^2(b_j)}{b_j(\omega - nq_j\Omega_{cj} - k_\parallel v_\parallel)}\frac{v_\perp}{v_\parallel}F_\parallel \tag{38}$$

$$F_\parallel = v_\parallel\left[\frac{\partial f_{0j}}{\partial v_\perp} + \frac{nq_j\omega_{cj}}{\omega}\left(\frac{v_z}{v_\perp}\frac{\partial f_{0j}}{\partial v_\perp} - \frac{\partial f_{0j}}{\partial v_\parallel}\right)\right] \tag{39}$$

The definitions for the plasma frequency ω_{pj}, the ion cyclotron frequency Ω_{cj} and the sign of the particle charge q_j have been used in (38) and (39). Additionally, the Bessel function of the first kind makes its appearance in (38), with argument $b_j = q_j k_\perp v_\perp/\Omega_{cj}$. When $\vec{B}_0$ is again assumed to be in the $\hat{z}$-direction, $k_\perp = \sqrt{k_x^2 + k_y^2}$ is the component of the wave vector perpendicular to $\vec{B}_0$, and $k_\parallel = k_z$ is the component of the wavevector parallel to $\vec{B}_0$.

This definition for the dielectric tensor can be used for any differentiable equilibrium distribution function f_{0j}. However, significant simplifications in the expressions can be made under the assumption that there are Maxwellian velocity distributions both perpendicular and parallel to the background magnetic field. Especially the integrals over velocity space simplify significantly. This choice was made

in this work as well, bringing the final expressions for the dielectric tensor that have been utilized in this work [44].

$$\vec{\vec{\epsilon}}_j = \begin{pmatrix} K_{1j} + \sin^2(\psi)\, K_{0j} & K_{2j} - \cos(\psi)\sin(\psi)\, K_{0j} & \cos(\psi)\, K_{4j} + \sin(\psi)\, K_{5j} \\ -K_{2j} - \cos(\psi)\sin(\psi)\, K_{0j} & K_{1j} + \cos^2(\psi)\, K_{0j} & \sin(\psi)\, K_{4j} - \cos(\psi)\, K_{5j} \\ \cos(\psi)\, K_{4j} - \sin(\psi)\, K_{5j} & \sin(\psi)\, K_{4j} + \cos(\psi)\, K_{5j} & K_{3j} \end{pmatrix} \tag{40}$$

$$\psi = \tan^{-1}\left(\frac{k_y}{k_x}\right) \tag{41}$$

$K_{0j} \dots K_{5j}$ are six similar expressions for contributions to this tensor. ψ denotes the angle that the perpendicular wavevector component makes with the $\hat{x}$-axis and can be determined using the four-quadrant arctangent.

Before introducing $K_{0j} \dots K_{5j}$, some dimensionless numbers need to be introduced. λ_j is the argument of the modified Bessel functions of the first kind $I_n(\lambda_j)$ of order n, and is related to ratio between the perpendicular wavelength and the Larmor radius. For the IBW, $\lambda_j \sim 1$. ζ_{nj} is the argument of the plasma dispersion function $Z(\zeta_{nj})$, and is related to the ratio between the antenna frequency and the n-th cyclotron frequency harmonic [45].

$$\lambda_j = \frac{1}{2} k_\perp^2 \rho_{Lj}^2 \tag{42}$$

$$\zeta_{nj} = \frac{\omega + n\Omega_{cj}}{k_\parallel v_j} = \frac{\omega/\Omega_{cj} + n}{k_\parallel \rho_{Lj}} \tag{43}$$

The Larmor radius ρ_{Lj} and the thermal velocity v_j are as defined in equations (3) and (4). The derivative of the modified Bessel function of order n with respect to λ_j is denoted by $I_n'(\lambda_j)$, and the derivative of the plasma dispersion function with respect to ζ_{nj} is denoted by $Z'(\zeta_{nj})$. The 6 components $K_{0j} \dots K_{5j}$ are now given by:

$$K_{0j} = 2\frac{\omega_{pj}^2 e^{-\lambda_j}}{\omega k_z v_j} \sum_{n=-\infty}^{\infty} \lambda_j (I_n - I_n') Z(\zeta_{nj}), \tag{44}$$

$$K_{1j} = \frac{\omega_{pj}^2 e^{-\lambda_j}}{\omega k_z v_j} \sum_{n=-\infty}^{\infty} \frac{n^2 I_n}{\lambda_j} Z(\zeta_{nj}), \tag{45}$$

$$K_{2j} = i\frac{q_j \omega_{pj}^2 e^{-\lambda_j}}{\omega k_z v_j} \sum_{n=-\infty}^{\infty} n(I_n - I_n') Z(\zeta_{nj}), \tag{46}$$

$$K_{3j} = -\frac{\omega_{pj}^2 e^{-\lambda_j}}{\omega k_z v_j} \sum_{n=-\infty}^{\infty} I_n \zeta_{nj} Z'(\zeta_{nj}), \tag{47}$$

$$K_{4j} = \frac{k_\perp \omega_{pj}^2 e^{-\lambda_j}}{2k_z \omega \Omega_{cj}} \sum_{n=-\infty}^{\infty} \frac{n I_n}{\lambda_j} Z'(\zeta_{nj}), \tag{48}$$

$$K_{5j} = i\frac{k_\perp q_j \omega_{pj}^2 e^{-\lambda_j}}{2k_z \omega \Omega_{cj}} \sum_{n=-\infty}^{\infty} (I_n - I_n') Z'(\zeta_{nj}). \tag{49}$$

In principle, an infinite number of terms should be added to $K_{0j} \dots K_{5j}$ to get the exact value, but in practice it suffices to limit the amount of terms in the sum, because the sum converges at a reasonable speed for low λ_j. For all hot-plasma simulations, n between $[-4,4]$ was used. Note that $K_{0j} \dots K_{5j}$ are position-dependent, because the density, magnetic field and temperature are not constant throughout the plasma.

Furthermore, special care must be taken when $k_\perp \to 0$, or $k_\parallel \to 0$. Numerical treatment of these points is given in Appendix D and Appendix E. Additionally, the product $e^{-\lambda_j} I_n(\lambda_j)$ has poor numerical properties for large λ_j.

2.6.2 The plasma current density in a hot plasma

The formulation of the dielectric tensor in a hot plasma (40) assumes that the background medium is homogeneous, such that the plasma current density in the spectral domain can be written as a simple multiplication between the conductivity tensor spectrum and the electric field spectrum.

$$\vec{J}_p(\vec{k}) = \vec{\vec{\sigma}}(\vec{k}) \cdot \vec{E}(\vec{k}) \tag{50}$$

The time-harmonic Helmholtz equation (11) however deals with electric fields in the spatial domain. Additionally, in a fusion reactor the background medium is inhomogeneous. In general, the response of the plasma to the electric field is non-local, and dispersive in both space and time. It can be expressed as [23] [46]:

$$\vec{J}_p(\vec{r}, t) = \int_{\infty}^{t} \int_{-\infty}^{\infty} \vec{\vec{\sigma}}(\vec{r}, \vec{r}', t, t') \cdot \vec{E}(\vec{r}', t') d\vec{r}' \, dt'. \tag{51}$$

The spatial dependence of $\vec{\vec{\sigma}}$ is caused by the spatially-dependent plasma properties like its density n and temperature T, and the equilibrium magnetic field $\vec{B}$ [23]. However, nearly all literature provides descriptions of the dielectric tensor in the spectral domain. Additionally, the convolution-like description of the plasma current in (51) is not desired in numerical applications.

Therefore, the assumption is often made that the plasma is locally homogeneous, which means that the Larmor radius is much smaller than the device size: $\rho_L \ll L$. Filling in the typical JET parameters reveals that the Larmor radius is at most 4mm in the core of the machine. The simulation domain on the other hand is 2 meters wide, which validates the assumption.

In this case, the spectrum of the plasma current at each test position x_t in the plasma is written using the spectrum of the electric field $\vec{E}(k_x)$ and the dielectric tensor spectrum, evaluated at x_t: $\vec{\vec{\sigma}}(k_x)|_{x=x_t}$. The inverse Fourier transform is used to obtain an expression in the spatial domain. Additionally, the assumption is made again that the electric and magnetic fields are time-harmonic.

$$\vec{J}_p(x) = \int_{\infty}^{\infty} \vec{\vec{\sigma}}(k_x)|_{x=x_t} \cdot \vec{E}(k_x) e^{ik_x x} dk_x \tag{52}$$

To allow for a slightly cleaner notation, $\vec{\vec{\sigma}}(k_x)|_{x=x_t}$ is denoted as $\vec{\vec{\sigma}}(x, k_x)$ from now on. Similar to the cold plasma case, an expression like (18) is used to obtain a dielectric tensor. Substituting the equation of the plasma current in the Helmholtz equation yields the integro-differential problem that is used in all three hot-plasma models.

$$-\nabla \times \nabla \times \vec{E}(x) + k_0^2 \int_{-\infty}^{\infty} \vec{\vec{\epsilon}}(x, k_x) \cdot \vec{E}(k_x) e^{ik_x x} dk_x = 0 \tag{53}$$

2.6.3 The hot plasma dispersion relation

At a similar method to the cold plasma, the dispersion relation for a hot plasma can be determined. For a homogeneous hot, magnetized plasma, there is a similar wave equation as for the cold plasma, equation (28).

$$\vec{k} \times \vec{k} \times \vec{E}_k + k_0^2 \, \vec{\vec{\epsilon}}(\vec{k}) \cdot \vec{E}_k = 0 \tag{54}$$

However, in contrast to the cold plasma, the hot plasma dielectric tensor is a function of $\vec{k}$ as well. Therefore, it is difficult to obtain analytic solutions for the dispersion relation, except for some simplified cases. Examples are propagation purely perpendicular to $\vec{B}_0$, such that $k_\parallel = 0$. For a general case however, the roots to (54) must be found by numerical methods. This section discusses this procedure and shows the results for the same JET case that was used in the cold plasma case. First of all, (54) can be written as a matrix system of the form:

$$\left(\begin{bmatrix} -(k_y^2 + k_z^2) & k_x k_y & k_x k_z \\ k_x k_y & -(k_x^2 + k_z^2) & k_y k_z \\ k_x k_z & k_y k_z & -(k_x^2 + k_y^2) \end{bmatrix} + k_0^2 \, \vec{\vec{\epsilon}}(\vec{k}) \right) \cdot \begin{bmatrix} E_x \\ E_y \\ E_z \end{bmatrix} = 0. \tag{55}$$

Solutions are searched for k_x in the complex plane; for k_y and k_z values are assumed based on typical values seen for the JET antenna.

A very basic rootfinder is implemented in order to search for solutions. The strategy consists of sampling every grid point at 2000 k_x positions, between $-10^6 < k_x < 10^6$. For each of the samples, the determinant is calculated numerically. At the point where the determinant shows a transition through 0, a solution is noted down. The results are plotted in the figure below.

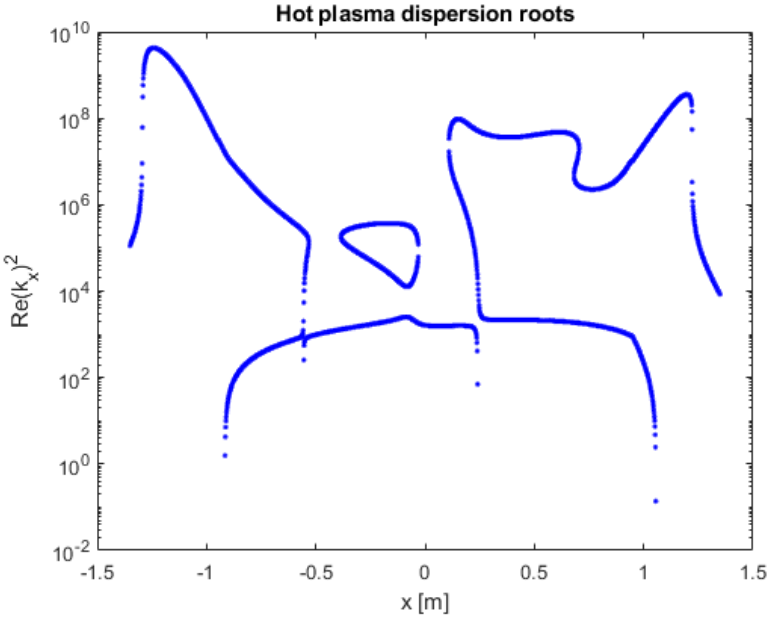

Figure 10: Some hot plasma dispersion roots. This graphs shows only a part of the possible waves in the plasma. More roots can be found by a more advanced rootfinder that searches for complex solutions.

If the objective is to find more roots, the algorithm must be adjusted. Right now, only real-valued k_x are investigated, but the cold plasma case already demonstrates that this is insufficient. As such, only weakly damped waves are found with this method: evanescent waves, or strongly damped waves are not discovered. Still, it is interesting so compare the results to Figure 7. In Figure 10 some high-wavenumber solutions are present which do not occur in the cold plasma. Even more, around the hydrogen ion cyclotron resonance layer the dispersion roots approach each other, which provides a way to couple

energy from the low-wavenumber wave, the fast wave, to the high-wavenumber wave, the Ion Bernstein Wave.

This coupling, or mode conversion, is a peculiar feature of the hot plasma simulations. The mode-converted IBW can clearly be seen in Figure 11. To the left of $x = 0$, part of the long-wavelength fast wave is converted into the short-wavelength ion Bernstein wave. This wave quickly damps out as its energy gets absorbed. This mode conversion process is not unique for fusion, it can also be observed in other magnetized plasmas, for example Earth's magnetosphere [31].

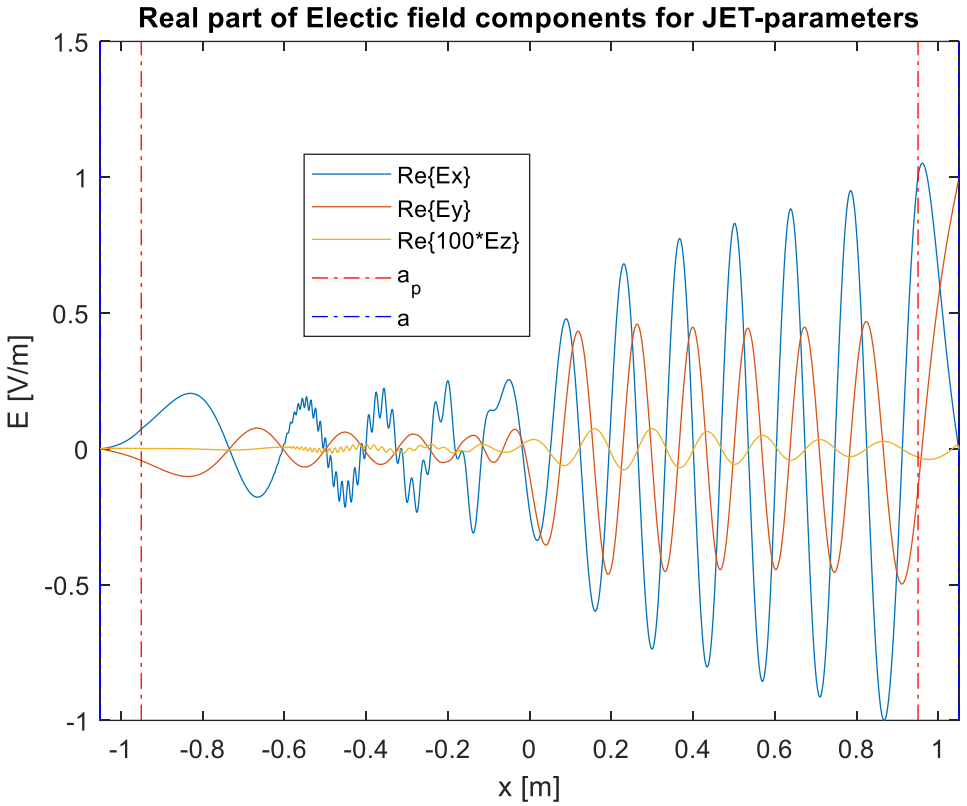

Figure 11: Hot plasma results for the same JET scenario as Figure 8. Note how in the hot plasma case energy from the long-wavelength fast wave is coupled into the short-wavelength Ion Bernstein Wave, a process referred to as mode-conversion.

3 Description of the one-dimensional plasma model

All models that have been created share a comparable setup of the simulation. This chapter discusses some details on the positioning of the 1D grid, the boundary conditions, the plasma profiles and the excitation. Additionally, the spill-over areas are discussed that are needed for the all-orders model.

3.1 Simulation domain

The propagation of electromagnetic waves in a fusion plasma is dependent on the background magnetic field. As the cold plasma and the hot plasma dielectric tensors point out, a magnetized plasma is an anisotropic medium. For correct simulations of the wave behavior, it is therefore necessary to consider the magnetic field topology. Due to the axisymmetry of tokamaks it often suffices to perform 2-dimensional wave simulations. The two-dimensional simulation domain would then be a cross section of the fusion reactor plasma. Stellarators have a more complex magnetic geometry, and as such they might benefit from three-dimensional simulations of the entire fusion reactor.

However, 3D and even 2D simulations of acceptable resolution are too demanding for a regular PC or laptop, hence all simulations in this work are executed on a 1-dimensional grid. A very simple torus geometry is assumed, with major radius R_0, and minor radius a [47]. Triangularity and elongation are not considered. The High Field Side (HFS), where the background magnetic field is highest, is in the center of the machine, while the Low Field Side (LFS) is on the outboard side of the machine. A rendering of a Tokamak, with a JET-like ratio of minor radius to major radius, is provided in Figure 12.

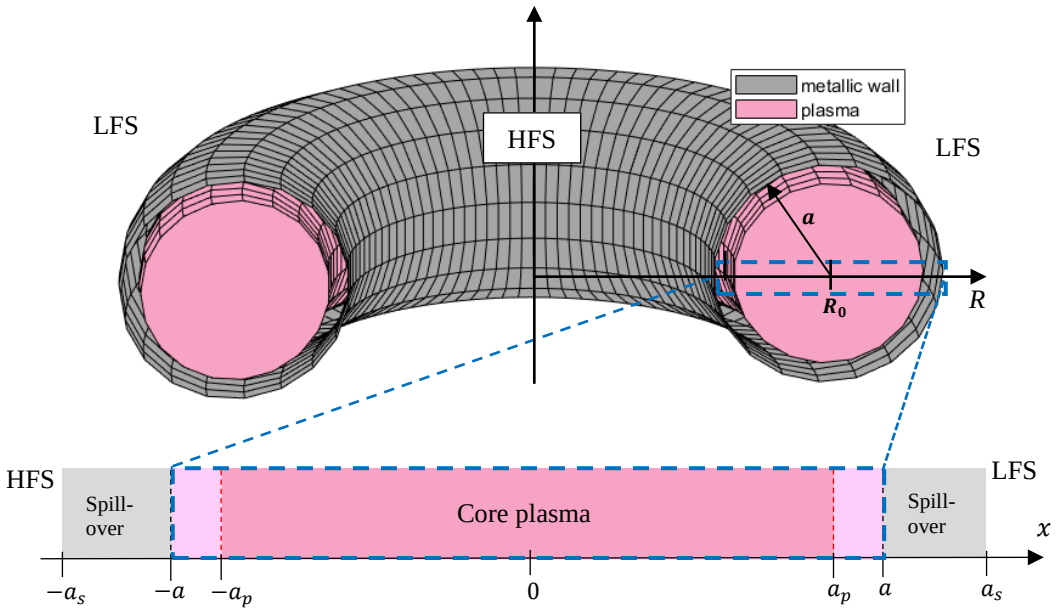

Figure 12: A simple rendering of a tokamak with a JET-like ratio ($R_0 = 2.97m, a_p = 0.95m$) (top), and the corresponding 1D simulation domain (bottom). Note that $x = -a$ and $x = a$ match the metallic walls of the machine, while the spill-over areas bear no physical correspondence.

In Figure 12, R is the radial coordinate, with $R = 0$ in the center of the machine, and $R = R_0$ corresponding to the major radius of the machine. In order to simplify matters a bit, the shifted radial coordinate x is introduced, such that $x = R - R_0$. The 1-dimensional physical domain runs from the wall on the HFS at $x = -a$, through the plasma core at $x = 0$, to the wall on the LFS at $x = a$. Excitation takes place on the wall on the LFS.

The boundaries of the full simulation domain are located at $x = \pm a_s = \pm L/2$. The region between a and a_s is a "spill-over" region, and bears no physical meaning. This is discussed in more detail in section 3.5. Finally, a_p is the transition point between the parabolic-like core plasma profiles, shown in dark pink, and the exponentially decaying edge plasma profiles, shown in light pink for the densities and temperatures.

3.2 Background magnetic field profile

Tokamaks like JET and ITER house an impressive set of large, powerful magnets which, together with a strong current through the plasma, generates the complex magnetic topology required to confine the plasma [48]. Especially in 2D and 3D simulations, this complex field topology should be accounted for in its entirety, as it influences the propagation of the waves. However, in this simplified model, only the toroidal magnetic field component, $\vec{B}_\theta$, is considered, as it is the strongest field component. The toroidal direction in a Tokamak is the direction "along the donut", as visualized in Figure 13.

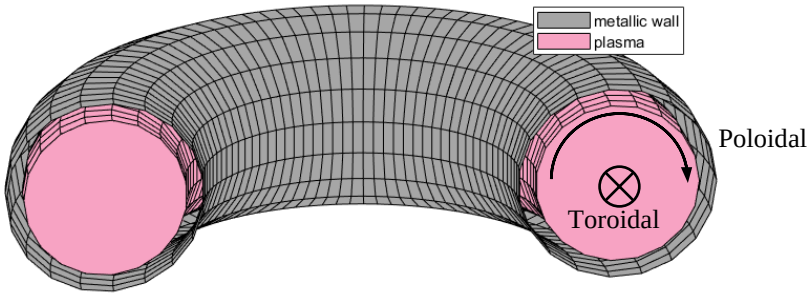

Figure 13: Sketch of the Toroidal direction and the poloidal direction within a Tokamak.

In the 1D geometry under consideration, the assumption is made that the background magnetic field is pointed into the paper, in the $\hat{z}$-direction, and that it is only a function of the shifted radial coordinate x, such that

$$\vec{B}(x) = \vec{B}_z(x) = B_z(x)\hat{z}. \tag{56}$$

The field strength $B_\theta(x)$ is strongest at the HFS, and lowest at the LFS. This is represented by a $1/R$ decay. The profiles are generated based on B_0, the magnetic field in the core of the plasma, at $x = 0$:

$$B_\theta(x) = \frac{B_0 R_0}{R} = B_0 \frac{R_0}{x + R_0}. \tag{57}$$

This profile has been visualized in Figure 14 for the JET parameters used throughout this work: $R_0 = 2.97\text{m}$ and $B_0 = 3.45\text{T}$.

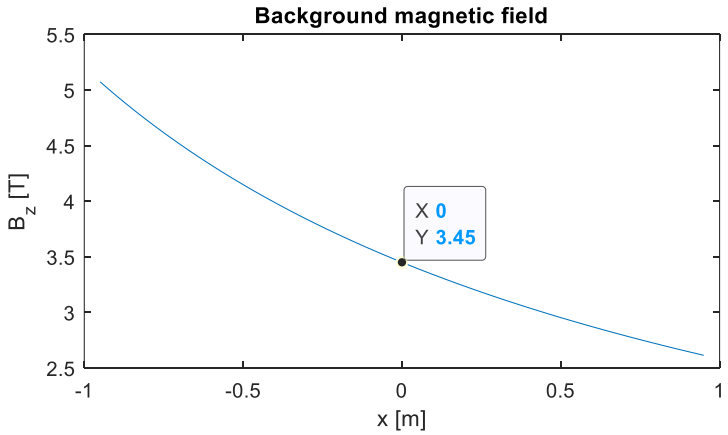

Figure 14: Background magnetic field strength profile $B_z(x)$ for JET parameters: $B_0 = 3.45T$ and $R_0 = 2.97m$. The field is highest at the HFS, and lowest at the LFS, and decays as $1/R$.

3.3 Temperature and Density profiles

The temperature and density profile, like the magnetic field profile, are simplified compared to reality. In the core plasma, both profiles are defined by the parabolic-like functions $G(x)$ described in the work of van Eester and Lerche [19]:

$$G_{core}(x) = (G_0 - G_s)\left(1 - \left(\frac{x}{a_p}\right)^2\right)^{\alpha} + G_s. \tag{58}$$

G_0 is the temperature or density on the magnetic axis, at $x = 0$. a_p is the transition point between the parabolic core plasma profiles and the exponentially decaying edge plasma profiles, and G_s is the density or temperature at this point. Finally, α is a shaping parameter. For the density profiles, $\alpha_n = 1$, yielding a parabola, while the temperatures have a stronger peaked nature, with $\alpha_T = 1.5$.

In the edge, the profiles decay exponentially, making sure they are continues with the core profiles at the transition point a_p.

$$G_{edge}(x) = G_s e^{-\frac{|x-a_p|}{\lambda}} \tag{59}$$

λ can be used to tune how fast the profile decays. Both for the temperature and the density $\lambda_T = \lambda_n = 0.05$ is chosen. The profiles are plotted in Figure 15.

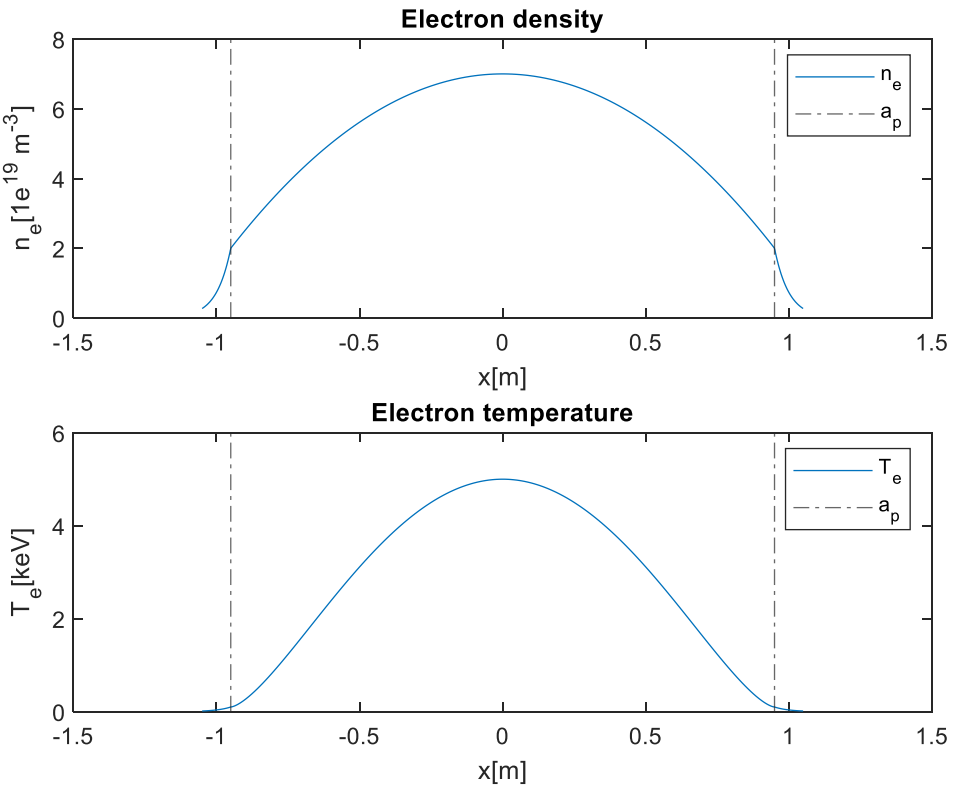

Figure 15: Electron temperature and density for JET parameters. The transition point between the parabolic-like core plasma profiles and the exponentially decaying edge plasma profiles is indicated with the dashed lines.

3.4 Boundary conditions

In this system, the boundary conditions are chosen such that the left wall, at $x = -a$, is a Perfect Electric Conductor (PEC). A PEC wall does not allow electric fields parallel to the wall, and as the walls are in the $y - z$ plane, E_y and E_z are forced to 0. At the right wall, at $x = a$, the excitation is modeled by setting a finite value for E_y. As discussed in section 2.4, this is comparable to the common situation where a Faraday screen blocks the electric field component parallel to the magnetic field, E_z, from passing through. This gives four Dirichlet boundary conditions for E_y and E_z:

$$E_y|_{x=-a} = 0, \tag{60}$$

$$E_z|_{x=-a} = 0, \tag{61}$$

$$E_y|_{x=a} = 1, \tag{62}$$

$$E_z|_{x=a} = 0. \tag{63}$$

These boundary conditions are added to the system matrix via Lagrange multipliers, instead of replacing existing equations [49] [50]. A total of 4 rows and 4 columns are added to the system matrix by conditions (60)-(63). All models except the all-orders model receive three extra conditions due to the periodicity requirements.

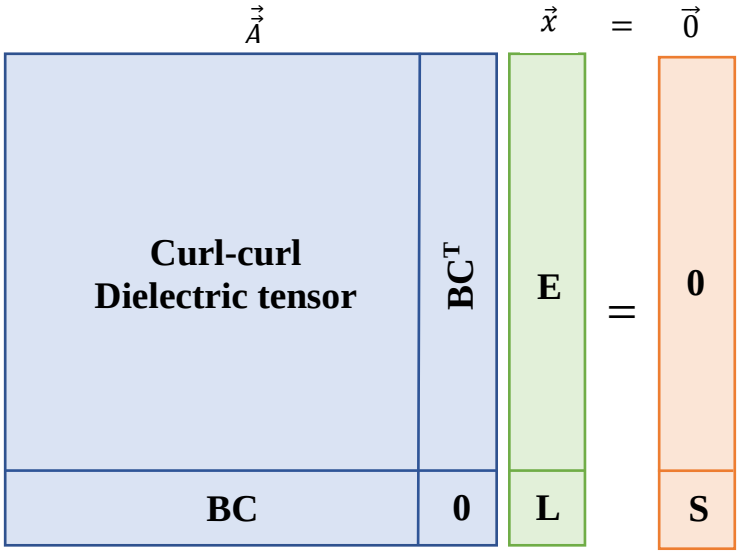

Figure 16: Structure of system matrix used for the simulations. The left-upper block contains the "actual" equations, linear in E. Amongst others it contains expressions for the curl-curl and the dielectric tensor contributions. The boundary conditions are added as Lagrange multipliers. (Recreated from [50])

After solving the linear system, the solution contains both the electric field values and the values for the Lagrange multipliers. Only the electric field values are kept, the Lagrange multipliers are discarded.

3.5 Spill-over areas

The Fourier-type approach from the all-orders model prefers fields that are periodic across the simulation domain. This is incompatible with the actual field structure. The one-sided excitation on E_y at the right wall for example results in a clear aperiodic situation, as the PEC boundary condition at the left wall forces E_y to 0. Additionally, there are evanescent regions in the plasma edge, which are typically difficult to represent with a Fourier series.

To deal with this problem, a small, non-physical, spill-over area is added on both sides of the simulation domain. In these regions, the electric fields can regain periodicity. To demonstrate how this works, the results for the all-orders model, discussed in section 5.1, is shown.

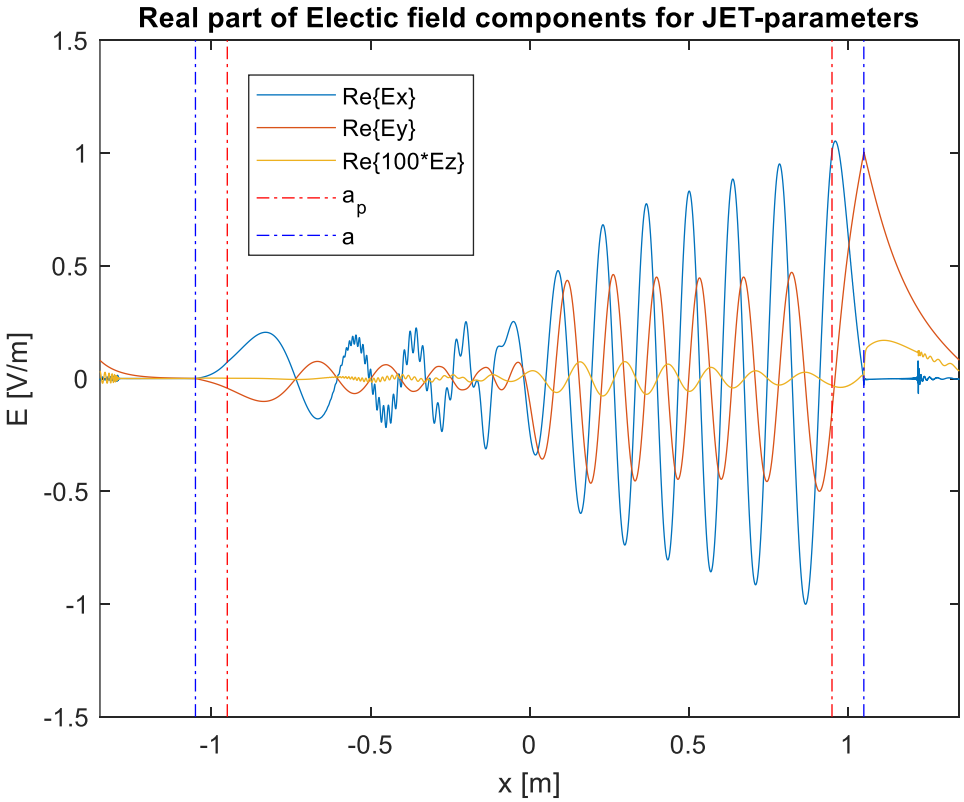

Figure 17: Electric field results from the all-orders model, demonstrating the spill-over areas for $|x| > a$. These areas do not bear any physical significance, and as such they are cut off in all plots to follow.

The region beyond $|x| > a$ does not bear any physical meaning, and its only purpose is to allow for the solutions to be periodic. In all further results these regions cut away to avoid confusion. The spill-over approach is used in all models, even for the purely differential models that do not necessarily require it, to keep the results as similar as possible.

3.6 Computational grid and Finite Difference scheme

For all models an equidistant grid is used with N grid points. The total simulation domain has a length L, spanning from $x = -a_s$ to $x = a_s$. The spacing h between the grid points is given by:

$$h = \frac{L}{N-1}. \tag{64}$$

The edges are not sampled twice, so the leftmost grid point is located at $x_1 = -a_s$, and the rightmost grid point is located at $x_N = a_s - h$. This avoids adding duplicate information to the matrix system, as the periodic boundary condition enforces continuity of the electric fields at $x = \pm a_s$. Neglecting to take this into account can result in ill-conditioned system matrices.

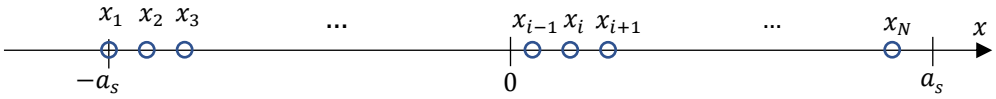

Figure 18: Sketch of the distribution of the gridpoints. The nodes are spaced at equal distance h from each other. The edge is not sampled twice: there is a sample at $x_1 = -a_s$ and a sample at $x_N = a_s - h$.

The cold plasma model, and two of the hot plasma models, use a Finite Difference (FD) scheme [51]. Other numerical schemes like Finite Element (FEM) might be more suitable for 2D or 3D simulations, or for non-uniform meshes. For a one-dimensional, equidistant grid however, the FD scheme is one of the easiest-to-implement numerical methods for discretizing derivatives. As an example, the second order derivative of the electric field with respect to x, at a discretized position x_i, can be written as:

$$\frac{\partial E(x)}{\partial x^2}\Big|_{x=x_i} \approx \frac{E_{i-1} - 2E_i + E_{i+1}}{h^2}. \tag{65}$$

Here E_{i-1}, E_i and E_{i+1} refer to the values of the electric fields at the nodes x_{i-1}, x_i and x_{i+1}. This is an example of a central difference scheme with second order accuracy. The coefficients for the finite difference scheme can be found in literature for many derivative orders, and for various orders of accuracy [52].

The cold plasma model (chapter 4) and the truncated Taylor series model (chapter 5.2) both use finite difference schemes of second order accuracy. These models only contain the first and second order derivative of $\vec{E}$ with respect to x. The truncated polynomial fit model (chapter 5.3) requires coefficients for much higher-order derivatives.

4 Cold plasma model

The cold plasma model is implemented based on the discussion of the theory in chapter 2.5, as a starting point for the more advanced hot-plasma simulations. The benefit of taking the cold plasma as a starting point, is that it will be easy to test the numerical schemes of the hot plasma models. In all hot plasma models that are created, the hot plasma dielectric tensor can be swapped for the cold plasma dielectric tensor by changing a single line in the code, to show that the numerical scheme can reproduce the cold plasma results. This chapter mainly discusses the implementation details, the results from this model can be seen in Figure 8.

4.1 Implementation of the Helmholtz equation

Even though the double curl of the electric field, $\nabla \times \nabla \times \vec{E}(\vec{r})$, might look difficult to implement, it is not the most challenging contribution to equation (11). A shorthand notation is used, where the curl-curl is written as the tensor $\overrightarrow{\overrightarrow{D}}$, acting on $\vec{E}$. When curvature of the magnetic field is not included, $\overrightarrow{\overrightarrow{D}}$ has a rather simple expression:

$$-\nabla \times \nabla \times \vec{E}(\vec{r}) = \overrightarrow{\overrightarrow{D}}\vec{E}(\vec{r}) = \begin{bmatrix} \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2} & -\frac{\partial^2}{\partial x \partial y} & -\frac{\partial^2}{\partial x \partial z} \\ -\frac{\partial^2}{\partial x \partial y} & \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial z^2} & -\frac{\partial^2}{\partial y \partial z} \\ -\frac{\partial^2}{\partial x \partial z} & -\frac{\partial^2}{\partial y \partial z} & \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} \end{bmatrix} \vec{E}(\vec{r}) \tag{66}$$

The expression for $\overrightarrow{\overrightarrow{D}}$ can be further simplified for the 1D case under consideration. The only free direction of the electric field $\vec{E}(\vec{r})$ is the x-direction; y- and z-dependence is captured by assuming a plane wave solution in $\hat{y}$- and $\hat{z}$-direction, with the wave numbers k_{y0} and k_{z0} respectively. Suitable choices for these wavenumbers are found by an a-priori assumption of the antenna excitation in these directions, either by experimental data, or through higher-dimensional models. Under this assumption, $\vec{E}(\vec{r})$ is given by:

$$\vec{E}(\vec{r}) = \vec{E}(x) e^{i(k_{y0} y + k_{z0} z)} \tag{67}$$

This choice simplifies the double curl matrix $\overrightarrow{\overrightarrow{D}}$ significantly, as the derivatives with respect to y and z reduce to multiplications with ik_y and ik_z respectively. Therefore, without curvature, the final expression for $\overrightarrow{\overrightarrow{D}}$ becomes:

$$\overrightarrow{\overrightarrow{D}} = \begin{bmatrix} -\left(k_y^2 + k_z^2\right) & -ik_y \frac{\partial}{\partial x} & -ik_z \frac{\partial}{\partial x} \\ -ik_y \frac{\partial}{\partial x} & -k_z^2 + \frac{\partial^2}{\partial x^2} & k_y k_z \\ -ik_z \frac{\partial}{\partial x} & k_y k_z & -k_y^2 + \frac{\partial^2}{\partial x^2} \end{bmatrix} \tag{68}$$

A derivation of $\overrightarrow{\overrightarrow{D}}$ in the curved field line case has been included in Appendix C. Some extra terms appear, but the approach is very similar. Both versions are implemented in code, however all results are generated with corrections for curvature taken into account.

4.2 Formulation of the cold plasma system of equations

Using the matrix style of formulation for the curl-curl that is discussed in the previous section, and using the cold plasma dielectric tensor as described in section 2.5.1, a very compact notation for the Helmholtz equation can be created. For simplicity, the Helmholtz equation for the cold plasma case in 1D is repeated here.

$$-\nabla \times \nabla \times \vec{E}(x) + k_0^2 \vec{\vec{\epsilon}}(x) \cdot \vec{E}(x) = 0 \tag{69}$$

A common approach for all models, including the hot plasma models, is grouping the contributions for each partial derivative order. The curl-curl for example can be split into three submatrices $\vec{\vec{D}}_0$, $\vec{\vec{D}}_1$ and $\vec{\vec{D}}_2$, where $\vec{\vec{D}}_0$, $\vec{\vec{D}}_1$ and $\vec{\vec{D}}_2$ now contain respectively the contributions for the 0th, 1st and 2nd derivatives of $\vec{E}(x)$ with respect to x.

$$\vec{\vec{D}} = \vec{\vec{D}}_0 + \vec{\vec{D}}_1 \frac{\partial}{\partial x} + \vec{\vec{D}}_2 \frac{\partial^2}{\partial x^2} \tag{70}$$

$$\vec{\vec{D}}_0 = \begin{bmatrix} -\left(k_y^2 + k_z^2\right) & 0 & 0 \\ 0 & -k_z^2 & k_y k_z \\ 0 & k_y k_z & -k_y^2 \end{bmatrix} \tag{71}$$

$$\vec{\vec{D}}_1 = \begin{bmatrix} 0 & -ik_y & -ik_z \\ -ik_y & 0 & 0 \\ -ik_z & 0 & 0 \end{bmatrix} \tag{72}$$

$$\vec{\vec{D}}_2 = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{73}$$

The dielectric tensor is already a 3×3 tensor, without need for derivatives with respect to x, and as such it can be combined with $\vec{\vec{D}}_0$. As such, it is possible to formulate the sub-matrices $\vec{\vec{M}}_0$, $\vec{\vec{M}}_1$ and $\vec{\vec{M}}_2$ that correspond to the contributions for the 0th, 1st and 2nd derivatives of $\vec{E}(x)$ respectively.

$$\vec{\vec{M}}_0 = \vec{\vec{D}}_0 + k_0^2 \vec{\vec{\epsilon}} \tag{74}$$

$$\vec{\vec{M}}_1 = \vec{\vec{D}}_1 \tag{75}$$

$$\vec{\vec{M}}_2 = \vec{\vec{D}}_2 \tag{76}$$

This combination forms the final form of the Helmholtz equation that is implemented in MATLAB.

$$\left(\vec{\vec{M}}_0 + \vec{\vec{M}}_1 \frac{\partial}{\partial x} + \vec{\vec{M}}_2 \frac{\partial^2}{\partial x^2}\right) \cdot \vec{E}(x) = 0 \tag{77}$$

The partial derivatives with respect to x are implemented using a finite difference scheme of second order accuracy. As such, each grid-point adds three 3×3 tensors to the system matrix, 27 contributions in total.

The results of this model for the JET parameters in appendix B have already been shown during the discussion of the cold plasma dispersion relation, Figure 8.

5 Hot plasma models

Three hot plasma models have been created, each with their own strengths and weaknesses. First, an "all-orders model" is discussed, which is accurate, but computationally very heavy. Secondly, a "truncated Taylor series model" is discussed, which is very fast, but fails to predict the damping of the Ion Bernstein Wave correctly. Finally, a "truncated polynomial fit model" is discussed, which aims to combine the best from both worlds.

5.1 All-orders model

The first model that is discussed, is an all-orders model. All-orders refers to the fact that no assumption is made on the value of λ_j. It has been created based on the All-Orders Spectral Algorithm (AORSA), from the work by Fred Jaeger [23]. This model was recreated for a 1-dimensional case, with a few minor changes. It is a computationally heavy, brute-force approach, and will be considered to be the benchmark throughout the report. A short summary of this model is given below, more details can be found in the AORSA paper.

5.1.1 Expression of the electric field as a sum of Fourier modii

The core of the AORSA code is the usage of the spectral collocation method, in which the electric field is expanded in a sum of Fourier modii. It is a simple, but effective way of discretizing the electric field spectrum, and it deserves a short description. The spectrum of the electric field, $\vec{E}(\vec{k})$, is related to the electric field in spatial coordinates, $\vec{E}(\vec{r})$, via the inverse spatial Fourier transform:

$$\vec{E}(\vec{r}) = \int \vec{E}(\vec{k}) e^{i\vec{k}\cdot\vec{r}} d\vec{k}. \tag{78}$$

The spectrum is a continuous, complex function of $\vec{k}$, and like the electric field in the spatial domain it is unknown. However, in a one-dimensional simulation domain of a certain length L, it is easy to depict field modii that fit an integer number of wavelengths in the cavity. The corresponding wavenumbers k_n are given by:

$$k_n = \frac{2\pi n}{L}. \tag{79}$$

Here n is an integer mode number, ranging from $-\frac{N}{2}$ to $\frac{N}{2}$. N is the number of gridpoints, but it is simultaneously the total number of modii. The expansion is only made in the $\hat{x}$-direction in the one-dimensional case. Similar to the cold plasma, a plane wave solution is assumed for the $\hat{y}$- and $\hat{z}$-direction, with wavenumbers k_{y0} and k_{z0} based on assumptions, measurements of simulations of the expected antenna spectrum in those directions. $\vec{E}(\vec{r})$ can then be expressed as the aforementioned sum of electric field modii. The goal of the all-orders simulation is finding the values for each electric field mode $\vec{E}_n$.

$$\vec{E}(x) = \sum_n \vec{E}_n e^{ik_n x} \tag{80}$$

5.1.2 Manipulation of the Helmholtz equation towards a system linear in $\vec{E}$

Using a similar expansion for the plasma current, an expression is obtained that is linear in $\vec{E}_n$. At each position in the plasma, the plasma current can be found based on the local spectrum of the hot plasma dielectric tensor [23],

$$\vec{J}_p(x) = \sum_n \sigma(x, k_n)\vec{E}_n e^{ik_n x}. \tag{81}$$

The expressions for the electric field and the plasma current can be substituted directly into the Helmholtz equation (11). Some rewrites, which are possible due to the distributive properties of the curl yields a system of equations in which $\vec{E}_n$ are the unknowns.

$$\sum_n \left[-\nabla \times \nabla \times \vec{E}_n + k_0^2 \vec{E}_n \vec{\vec{\epsilon}}(x, k_n)\right] e^{ik_n x} = 0 \tag{82}$$

This equation can be cleaned up using similar techniques as for the cold plasma model. First of all, the curl-curl is written in tensor form. It differs slightly from the cold plasma tensor in equation (68), in that the partial derivatives with respect to x are now replaced by ik_n.

$$\vec{\vec{D}} = \begin{bmatrix} -\left(k_{y0}^2 + k_{z0}^2\right) & k_n k_{y0} & k_n k_{z0} \\ k_n k_{y0} & -(k_n^2 + k_{z0}^2) & k_{y0} k_{z0} \\ k_n k_{z0} & k_{y0} k_{z0} & -\left(k_n^2 + k_{y0}^2\right) \end{bmatrix} \tag{83}$$

The curl-curl tensor $\vec{\vec{D}}$ can again be grouped together with the dielectric tensor $\vec{\vec{\epsilon}}$ to produce the combined tensor $\vec{\vec{M}}_n$,

$$\vec{\vec{M}}_n = \vec{\vec{D}}(k_n) + k_0^2 \vec{\vec{\epsilon}}(x, k_n), \tag{84}$$

which allows for a very compact notation of equation (82):

$$\sum_n \vec{\vec{M}}_n(x, k_n) e^{ik_n x} \vec{E}_n = 0. \tag{85}$$

Even though this entire section has been tailored to 1D simulations, extension to 2D and 3D is possible with minimal effort. In the AORSA code for example, a 2-dimensional derivation is performed, starting from an expansion in Fourier harmonics in both $\hat{x}$ and $\hat{y}$-direction.

5.1.3 Introduction of the weak form by using the Fast Fourier Transform

So far, the approach is near-identical to the approach taken in the AORSA code. At this point, the approach slightly differs. The system matrices that are obtained with the AORSA approach are typically nearly stochastic. A more regular matrix, with the dominant contributions on the main diagonal, is produced by writing down the weak formulation of (85) using test functions G_n.

$$G_n(x) = e^{-ik'_n \cdot x} \tag{86}$$

The weak form requires multiplication by G_n and integration over the entire domain.

$$\int_0^L e^{-ik'_n x} \left(\sum_n \vec{\vec{M}}_n(x, k_n) e^{ik_n x} \vec{E}_n \right) dx = 0 \tag{87}$$

Swapping the sum and integral, and regrouping components, gives the new expression for the Helmholtz equation.

$$\sum_n \left(\int_0^L e^{-ik'_n x} \vec{\vec{M}}_n(x, k_n) e^{ik_n x} dx \right) \vec{E}_n = 0 \tag{88}$$

The multiplication with the basis functions, and the integration, are computationally expensive: tests showed out the time required scaled approximately with N^3 for this formulation. Fortunately, due to the nature of the basis functions, this operation can be very efficiently executed by using a Fast Fourier Transform (FFT), bringing it down to a complexity of roughly $N \log N$. Some rewriting is required in order to shape (88) into a form that is compatible with MATLAB's Fourier transform. The definition of its built-in fft function is given by [53]:

$$Y(k) = fft(X) = \sum_{j=1}^{n} X(j) e^{-\frac{2\pi i}{N}(j-1)(k-1)} \tag{89}$$

First, the integral needs to be discretized, with h the grid spacing: $h = L/(N-1)$.

$$\sum_n \left(\sum_s e^{-ik'_n x_s} \vec{\vec{M}}_n(x_s, k_n) e^{ik_n x_s} \right) h \vec{E}_n = 0 \tag{90}$$

Here k'_n is a symmetric spectrum with N modes. n' ranges from $1 \dots N$, and the corresponding modes range from $-N_k \dots N_k$ $(N = 2N_k + 1)$. Furthermore, with $n = 1 \dots N$ as well, the expressions for k'_n are similar, given by:

$$k'_n = \frac{2\pi(n' - N_k - 1)}{L}, \tag{91}$$

$$k_n = \frac{2\pi(n - N_k - 1)}{L}. \tag{92}$$

The grid is sampled at "fft-compatible" gridpoints x_s, with $s = 1 \dots N$:

$$x_s = (s-1)\frac{L}{N}. \tag{93}$$

These expressions can be substituted in (90), yielding an expression that closely resembles MATLAB's fft transform.

$$\sum_n \left(\sum_{s=1}^{N} \left[e^{i(k_n+k_N)x_s} \vec{\vec{M}}_n(x_s,k_n) \right] e^{-\frac{2\pi i}{N}(s-1)(n'-1)} \right) h\vec{E}_n = 0 \tag{94}$$

Which yields effectively:

$$\sum_n \vec{\vec{M}}_n^{fft}(x_s,k_n)\vec{E}_n = 0, \tag{95}$$

$$\vec{\vec{M}}_n^{fft}(x_s,k_n) = \text{fft}\left[h e^{i(k_n+k_N)x_s} \vec{\vec{M}}_n(x_s,k_n) \right]. \tag{96}$$

As mentioned before, the power of the fast Fourier transform dramatically speeds up the computation of the integrals, changing it from a computational complexity of about N^3 to a complexity of roughly $N \log N$. With this reformulation, the matrix inversion is again the computational bottleneck.

5.1.4 Implementation in MATLAB

Except for the boundary conditions, the system matrix for the all-orders model has a fill factor of 100%. Each grid point adds N 3×3 tensors to the matrix. Therefore, there are a total of $3N \times 3N$ contributions in the core of the system matrix in a one-dimensional case. The boundary conditions add a few extra rows and columns, but this is a minor addition to the large, full core matrix. A sketch of the core matrix, as it is implemented in MATLAB, is provided in equation (97).

$$\begin{bmatrix} \vec{\vec{M}}_1^{fft}(x_1,k_1) & \vec{\vec{M}}_2^{fft}(x_1,k_2) & \dots & \vec{\vec{M}}_N^{fft}(x_1,k_N) \\ \vec{\vec{M}}_1^{fft}(x_2,k_1) & \vec{\vec{M}}_2^{fft}(x_2,k_2) & & \vdots \\ \vdots & & \ddots & \\ \vec{\vec{M}}_1^{fft}(x_N,k_1) & \dots & & \vec{\vec{M}}_N^{fft}(x_s,k_N) \end{bmatrix} \begin{bmatrix} \vec{E}_1 \\ \vec{E}_2 \\ \vdots \\ \vec{E}_N \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix} \tag{97}$$

Due to the expansion of the electric field in its spectral components, the Dirichlet boundary conditions for the electric field require some extra attention for the all orders model. The desired values for the electric field at the left and right wall are the same as those discussed in (60)-(63). The excitation at the right wall is taken as an example.

$$\sum_n E_{y,n} e^{ik_n a} = 1 \tag{98}$$

These boundary conditions are appended to the system matrix, using the Lagrange multiplier approach that is discussed in chapter 3.4, adding a total of 4 rows and 4 columns.

5.1.5 All-orders model results

A high-resolution simulation is presented, consisting of 5001 grid points/modii. This is about the limit that a PC with 16GB RAM can comfortably run before running into memory issues. A JET-like scenario is chosen, with the parameters as discussed in Appendix B, and profiles as discussed in chapters 3.2 and 3.3. It is a deuterium plasma, with a 5% hydrogen minority. The results from this all-orders simulation are used as a point of reference for all other hot-plasma simulations.

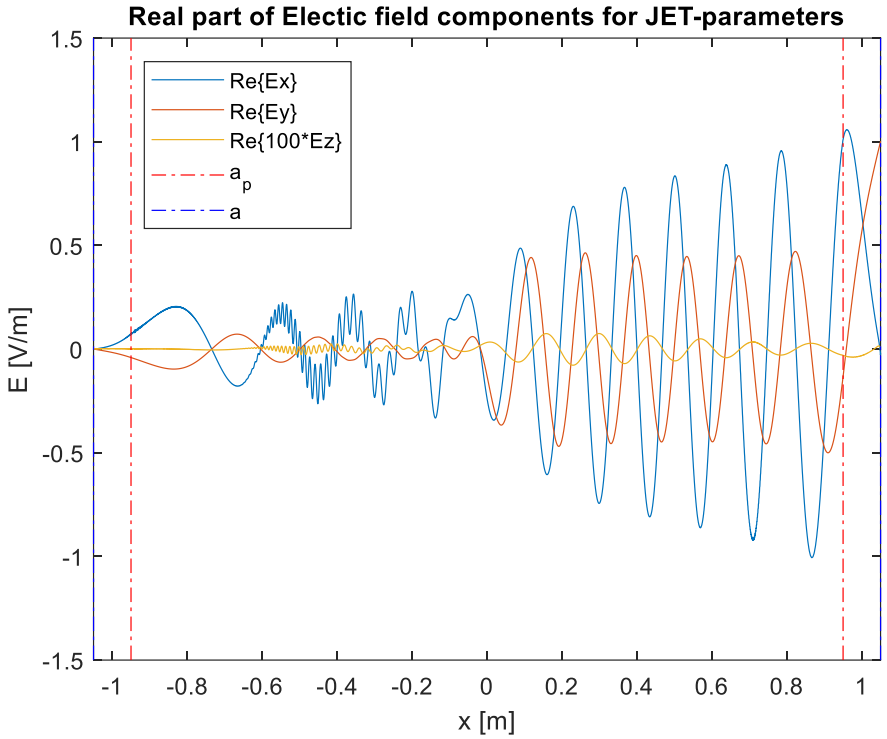

Figure 19: Real part of the electric field components, as calculated by the all-orders model. Only the region between $-a$ and a is plotted. The spill-over regions are cut away. Furthermore, the transition between the core plasma profiles and the edge plasma profiles is indicated with the dashed red lines. Note how the short-wavelength Ion Bernstein Wave is present in this hot-plasma case, in contrast to the cold plasma.

The fundamental hydrogen ion cyclotron resonance layer is present slightly to the right of $x = 0$. Around this point, the fast wave loses much of its amplitude, indicating absorption. In contrast to the cold-plasma results, a short-wavelength wave is present to the left of $x = 0$. This is the mode-converted Ion Bernstein Wave (IBW). The wave generated in the core of the machine, near the fundamental hydrogen cyclotron resonance. Visual inspection learns that the IBW is damped out by the time it reaches $x = -0.66\text{m}$.

5.2 Truncated Taylor Series model

Instead of following the approach taken by the all-orders model, where the entire equation is solved in the spectral domain, everything is now written in the spatial domain. The description that follows in this chapter shows just one of the possible routes that can be used. It was suggested by Dirk van Eester and adapted for 1D application. This model will make assumptions on the value of λ_j, which is incorrect for mode-converted Bernstein waves. Still, the approach taken is very elegant, and it lays the foundation for the truncated polynomial model, which is discussed in chapter 5.3.

5.2.1 Taylor Expansion of the dielectric tensor

As discussed before, an a-priori choice for the wavenumbers k_{y0} and k_{z0} are chosen in $\hat{y}$-, and $\hat{z}$-direction respectively, based on the antenna parameters, 2D or 3D simulations, or experimental results. Only k_x corresponds to a free direction. These three wavenumbers are grouped together in what will be called the continuous wave vector $\vec{k}_c$:

$$\vec{k}_c = [k_x \quad k_{y0} \quad k_{z0}\,]^T. \tag{99}$$

In this case, $\vec{\epsilon}(x,\vec{k})$ becomes only dependent on the local plasma parameters at a certain position x, and the wavenumber in $\hat{x}$-direction, k_x, via $\vec{k}_c$:

$$\vec{\epsilon}(x,\vec{k}) = \vec{\epsilon}(x,\vec{k}_c)\delta(k_y - k_{y0})\delta(k_z - k_{z0}). \tag{100}$$

The next step is writing down the Taylor expansion of $\vec{\epsilon}(x,k_x)$ around a test wavenumber k_{xt}, which together with the chosen k_{y0} and k_{z0} results in the test wavevector $\vec{k}_t$:

$$\vec{k}_t = [k_{xt} \quad k_{y0} \quad k_{z0}\,]^T. \tag{101}$$

The second order Taylor expansion is than simply given by:

$$\vec{\epsilon}(x,\vec{k}_c) \approx \vec{\epsilon}(x,\vec{k}_t) + (k_x - k_{xt})\vec{\epsilon}_{k_x} + \frac{1}{2}(k_x - k_{xt})^2\vec{\epsilon}_{k_xk_x}. \tag{102}$$

Where the compact notation $\vec{\epsilon}_{k_x}$ and $\vec{\epsilon}_{k_xk_x}$ denote respectively the first and second derivatives of $\vec{\epsilon}(x,\vec{k}_c)$ with respect to k_x, evaluated at $\vec{k}_t$:

$$\vec{\epsilon}_{k_x} = \left[\frac{\partial}{\partial k_x}\vec{\epsilon}(x,\vec{k}_c)\right]_{\vec{k}_c=\vec{k}_t}, \tag{103}$$

$$\vec{\epsilon}_{k_xk_x} = \left[\frac{\partial^2}{\partial k_x^2}\vec{\epsilon}(x,\vec{k}_c)\right]_{\vec{k}_c=\vec{k}_t}. \tag{104}$$

This approximation holds under the condition that the spectrum is narrow, such that a Taylor expansion around a certain $\vec{k}_t$ is justified. Note also that the derivatives act on each of the dielectric components separately, such that, for example:

$$\vec{\vec{\epsilon}}_{k_x} = \begin{bmatrix} \frac{\partial}{\partial k_x}\epsilon_{11}(x,\vec{k}_c) & \frac{\partial}{\partial k_x}\epsilon_{12}(x,\vec{k}_c) & \frac{\partial}{\partial k_x}\epsilon_{13}(x,\vec{k}_c) \\ \frac{\partial}{\partial k_x}\epsilon_{21}(x,\vec{k}_c) & \frac{\partial}{\partial k_x}\epsilon_{22}(x,\vec{k}_c) & \frac{\partial}{\partial k_x}\epsilon_{23}(x,\vec{k}_c) \\ \frac{\partial}{\partial k_x}\epsilon_{31}(x,\vec{k}_c) & \frac{\partial}{\partial k_x}\epsilon_{32}(x,\vec{k}_c) & \frac{\partial}{\partial k_x}\epsilon_{33}(x,\vec{k}_c) \end{bmatrix}_{\vec{k}_c=\vec{k}_t} . \quad (105)$$

A graph of $\epsilon_{11}(x,\vec{k}_c)$ is provided, together with its Taylor expansion around $k_{xt} = 0$ in Figure 20.

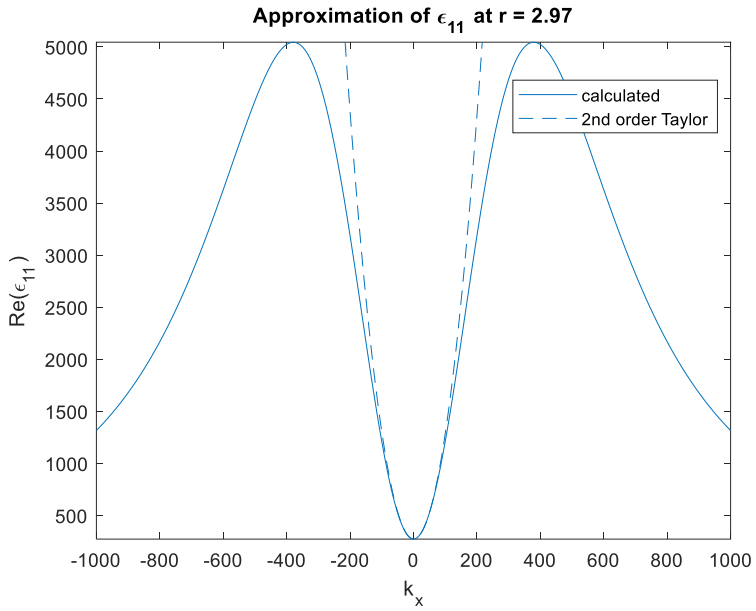

Figure 20: Comparison between the calculated spectrum of the dielectric tensor component ϵ_{11} (solid line) and the second order Taylor expansion (dashed line) around $k_{xt} = 0$. The spectrum is taken in the core of JET, under the same JET parameters used throughout the report. Only the real part of ϵ_{11} is plotted.

5.2.2 Differential operators and inverse Fourier transform

A very neat property of the inverse spatial Fourier transform is now used, which is the key to both the truncated Taylor model, and the truncated polynomial model. For an arbitrary function $f(k)$, the inverse Fourier of $k^n f(k)$ results in an n^{th} derivative of $f(x)$ [54]:

$$\int_{-\infty}^{\infty} (ik)^n f(k) e^{ikx} dk = \frac{d^n}{dx^n} f(x). \quad (106)$$

Note that this transform requires a bilateral Fourier transform: it assumed the full spectrum, from $-\infty$ to ∞ is considered. This might be a source for error, as the spectrum of $\vec{\vec{\epsilon}}$ is in principle infinitely wide, and both the truncated Taylor model and the truncated polynomial model fail to account for the full spectrum.

This property is quite important, as it allows writing the integro-differential wave equation from the hot plasma model as a purely differential model. Sparse, banded matrices are obtained, instead of the completely filled matrices, common to the all-orders model. Both the truncated Taylor series model, and the Truncated polynomial fit model, discussed in section 5.3, use this property.

5.2.3 Substitution in wave equation

In order to make full use of the inverse Fourier transform property described above, it is useful to rewrite the Taylor expansion, and group the contributions by powers of k_x.

$$\vec{\vec{\epsilon}}(x,\vec{k}_c) \approx \left(\vec{\vec{\epsilon}}(x,\vec{k}_t) - k_{xt}\vec{\vec{\epsilon}}_{k_x} + \frac{1}{2}k_{xt}^2\vec{\vec{\epsilon}}_{k_xk_x}\right) + \left(\vec{\vec{\epsilon}}_{k_x} - k_{xt}\vec{\vec{\epsilon}}_{k_xk_x}\right)k_x + \left(\frac{1}{2}\vec{\vec{\epsilon}}_{k_xk_x}\right)k_x^2 \tag{107}$$

The three groups that correspond to k_x^0, k_x^1 and k_x^2 will be denoted as the 3×3 tensors $\vec{\vec{\mathcal{E}}}_0$, $\vec{\vec{\mathcal{E}}}_1$ and $\vec{\vec{\mathcal{E}}}_2$ respectively:

$$\vec{\vec{\mathcal{E}}}_0 = \vec{\vec{\epsilon}}(x,\vec{k}_t) - k_{xt}\vec{\vec{\epsilon}}_{k_x} + \frac{1}{2}k_{xt}^2\vec{\vec{\epsilon}}_{k_xk_x}, \tag{108}$$

$$\vec{\vec{\mathcal{E}}}_1 = \vec{\vec{\epsilon}}_{k_x} - k_{xt}\vec{\vec{\epsilon}}_{k_xk_x}, \tag{109}$$

$$\vec{\vec{\mathcal{E}}}_2 = \frac{1}{2}\vec{\vec{\epsilon}}_{k_xk_x}. \tag{110}$$

This choice allows for a very compact notation of $\vec{\vec{\epsilon}}(x,\vec{k}_c)$:

$$\vec{\vec{\epsilon}}(x,\vec{k}_c) \approx \vec{\vec{\mathcal{E}}}_0 + \vec{\vec{\mathcal{E}}}_1k_x + \vec{\vec{\mathcal{E}}}_2k_x^2. \tag{111}$$

When the approximation for the dielectric tensor is then substituted in the one-dimensional hot-plasma wave equation, the intermediate result is:

$$-\nabla\times\nabla\times\vec{E}(x) + k_0^2\int\left[\vec{\vec{\mathcal{E}}}_0 + \vec{\vec{\mathcal{E}}}_1k_x + \vec{\vec{\mathcal{E}}}_2k_x^2\right]\cdot\vec{E}(k_x)e^{ik_xx}dk_x = 0. \tag{112}$$

At this point, the inverse Fourier transform property can be used, which transforms the powers of k_x into derivatives with of $\vec{E}(x)$ respect to x. It is at this point that the integro-differential equation becomes a purely differential equation, which can be solved much faster.

$$-\nabla\times\nabla\times\vec{E}(x) + k_0^2\left[\vec{\vec{\mathcal{E}}}_0 - i\vec{\vec{\mathcal{E}}}_1\frac{\partial}{\partial x} - \vec{\vec{\mathcal{E}}}_2\frac{\partial^2}{\partial x^2}\right]\vec{E}(x) = 0 \tag{113}$$

Finally, by writing the curl-curl again as the tensor components $\vec{\vec{D}}_0$, $\vec{\vec{D}}_1$ and $\vec{\vec{D}}_2$, the final expression in terms of the 3×3 tensors $\vec{\vec{M}}_0$, $\vec{\vec{M}}_1$ and $\vec{\vec{M}}_2$ is given, where the subscripts denote the contributions for the 0th, 1st and 2nd order derivatives respectively:

$$\left[\vec{\vec{M}}_0 + \vec{\vec{M}}_1\frac{\partial}{\partial x} + \vec{\vec{M}}_2\frac{\partial^2}{\partial x^2}\right]\vec{E}(x) = 0, \tag{114}$$

$$\vec{M}_0 = \vec{\vec{D}}_0 + k_0^2\vec{\vec{\mathcal{E}}}_0, \tag{115}$$

$$\vec{M}_1 = \vec{\vec{D}}_0 - ik_0^2\vec{\vec{\mathcal{E}}}_1, \tag{116}$$

$$\vec{M}_2 = \vec{\vec{D}}_2 - k_0^2\vec{\vec{\mathcal{E}}}_2. \tag{117}$$

5.2.4 Truncated Taylor Series model results

The results for the truncated Taylor series model is given in Figure 21. All parameters and profiles are kept the same as for the all-orders model: it is again a JET-like scenario, with 5% hydrogen in deuterium. Instead of 5001 grid points, 100.000 grid points were chosen, to highlight the strength, but also weakness of this approach.

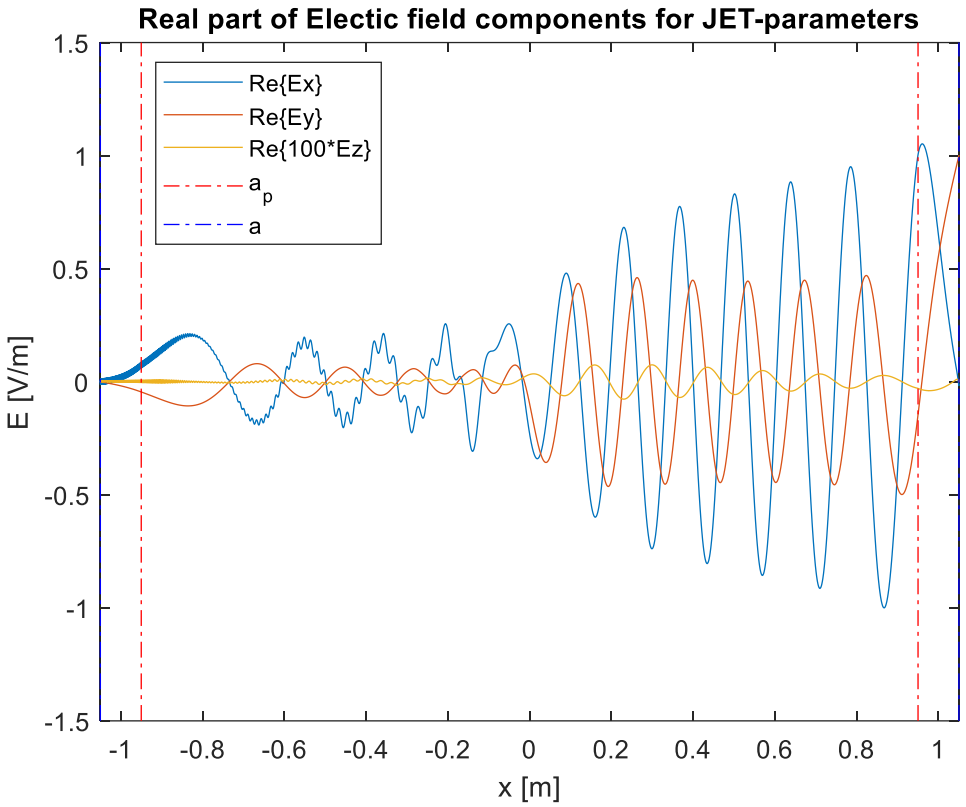

Figure 21: Real part of the electric field components, as calculated by the truncated Taylor series model. Only the region between $-a$ and a is plotted, bounded by the dashed blue lines, the spill-over regions are cut away. Furthermore, the transition between the core plasma profiles and the edge plasma profiles is indicated with the dashed red lines.

The results from this model mainly differs from the all-orders results in Figure 19, left of the cyclotron resonance layer. The damping of the IBW is much smaller, resulting in propagation of this wave until the left wall. Due to the decreasing temperature near the edge, the Larmor radius decreases and the IBW shrinks with it. The high resolution reveals that the IBW keeps propagating until it hits the left wall, at which point it has a wavelength of only 0.7mm. This in contrast to the all-orders model, where the wave doesn't even come close to leaving the plasma core. To summarize: this approach is very fast, even with 100.000 grid points, but damping of the IBW is underestimated. This is visible in the absorbed power as well, as will be discussed in chapter 6.2.

5.3 Truncated Polynomial Fit model

This model is the main product of this work. It is based on the observation that the dielectric tensor is, under most circumstances, a very smooth and slowly varying function of $k_\perp$, and this property is exploited by fitting an N_p^{th} order polynomial through the dielectric tensor components. A neat property of the Fourier transform than allows a reduction from an integro-differential model to a N_p^{th} order differential model. This offers significant benefits regarding computational complexity compared to the all-orders model, while not sacrificing too much accuracy, as will be shown in the results.

5.3.1 Approximating the dielectric tensor in $\vec{k}$-space

The dielectric tensor $\vec{\vec{\epsilon}}(\vec{r},\vec{k})$ is dependent on the local plasma conditions like magnetic field, temperature and densities, and the wavenumber of interest. In a 1D simulation, k_y and k_z at position x_i are known and fixed as well. Therefore, the 9 dielectric tensor components ϵ_{ij} at position x_i can be written as a function of k_x:

$$\epsilon_{ij}(\vec{r},\vec{k})|_{x=x_i} = \epsilon_{ij}(k_x). \tag{118}$$

ICRH simulations with AORSA suggest that ion Bernstein waves with wavelengths comparable to the Larmor radius are born: $k_x\rho_L \sim 1$ [23]. However, upon inspection of the dielectric tensor components it becomes clear that it should be possible to approximate the components up to $k_x\rho_L \sim 2$ with a polynomial of order N_p:

$$\epsilon_{ij}(x,k_x) \approx \sum_{p=0}^{N_p} c_{ij,p}(x)k_x^p. \tag{119}$$

This is visualized in Figure 22, where an 8-th order polynomial is fitted through $-1 < k_x\rho_L < 1$ for JET-like parameters around the hydrogen ion resonance layer, with Maxwellian distribution functions.

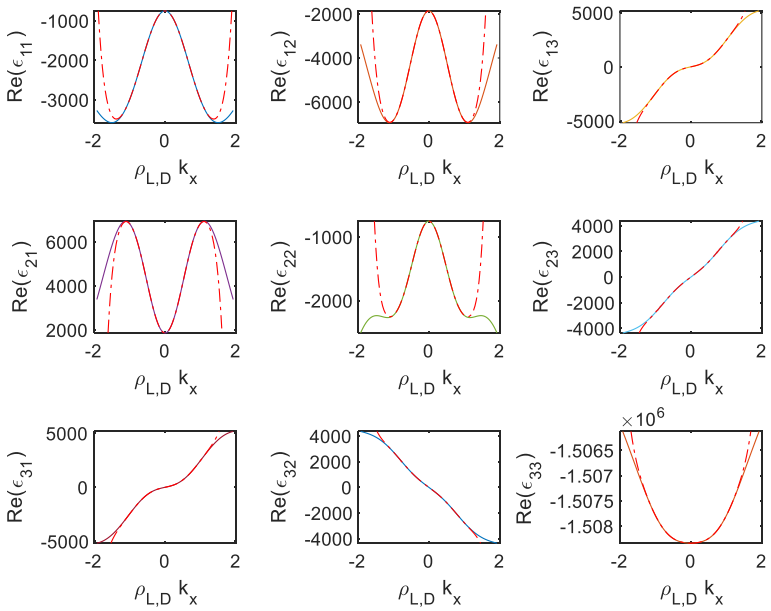

Figure 22: Demonstration of an 8-th order fit through $-1 < \rho_L k_x < 1$, for JET-like parameters around the hydrogen resonance layer. The fits are the dashed red lines, the computed values are the solid lines.

This expansion again allows usage of (106), the property of the inverse Fourier transform, that states that

$$\int_{-\infty}^{\infty} (ik)^n f(k) e^{ikx} dk = \frac{d^n}{dx^n} f(x). \tag{120}$$

Therefore, the inverse Fourier transform for the hot plasma current density can be translated to a sum of derivatives, up to order N_p. Per grid point x_i 9 fits are performed, one for each dielectric tensor component ϵ_{ij}. The coefficients for each power p of k_x are housed in the 3×3 tensor $\vec{\vec{c}}_p$, yielding:

$$\int \vec{\vec{\epsilon}}(x, k_x) \cdot \vec{E}(k_x) e^{ik_x x} dk_x \approx \sum_{p=0}^{N_p} -i^p \vec{\vec{c}}_p(x) \frac{d^p}{dx^p} \vec{E}(x). \tag{121}$$

Rewriting the curl-curl as the operator $\vec{\vec{D}}$, as in equation (66), and inserting (121) in the general Helmholtz equation , yields the intermediate differential form of the Helmholtz equation:

$$\left(\vec{\vec{D}}(x) + k_0^2 \sum_{p=0}^{N_p} -i^p \vec{\vec{c}}_p(x) \frac{d^p}{dx^p} \right) \vec{E}(x) = 0. \tag{122}$$

A very clean expression remains for the Helmholtz equation:

$$\sum_{p=0}^{N_p} \vec{\vec{M}}_p \frac{d^p}{dx^p} \vec{E}(x) = 0, \tag{123}$$

with $\vec{\vec{M}}_p$ the 3×3 tensor containing the contributions from $\vec{\vec{D}}_p$ and $\vec{\vec{c}}_p$. It should be clear that $\vec{\vec{D}}_p$ only delivers contributions for $p = 0,1,2$, while $\vec{\vec{c}}_p$ delivers contributions up to $p = N_p$.

$$\vec{\vec{M}}_p = \vec{\vec{D}}_p - k_0^2 i^p \vec{\vec{c}}_p(x) \tag{124}$$

5.3.2 Implementation in MATLAB

Equation (123) is implemented in MATLAB, using the expressions for the dielectric tensor, and the limit for $k_\perp \to 0$ is implemented as explained in Appendix E.

For every position in the plasma, the Larmor radius ρ_L of the lightest ion species in the plasma is determined. The window over which the fit is performed, ranging from $-k_{xmax} \dots k_{xmax}$, is scaled with the local Larmor radius. Effectively, this means that for the lower temperatures near the edge, a wider spectrum is considered for the fit. A "window size" parameter ζ is introduced for this purpose, such that the quantity $k_x \rho_{Lj}$ stays constant across the simulation domain. Note that ζ bears a lot of similarities with λ_j, equation (42).

$$k_{xmax} \rho_{Lmin} = \zeta \tag{125}$$

The dielectric tensor is then sampled at a minimum of $N_p + 1$ points in the range $-k_{xmax} \dots k_{xmax}$. A least squares fit is than performed through these points to find $N_p + 1$ coefficient matrices $\vec{\vec{c}}_p$. The contributions from the curl-curl are then added according to (124) to find $N_p + 1$ matrices $\vec{\vec{M}}_p$. A central difference scheme with 2nd order accuracy is used to find a_n^p, the coefficients for derivative order p at gridpoint n.

$$\vec{\vec{A}}_n = \sum_{p=0}^{N_p} \frac{a_n^p}{h^p} \vec{\vec{M}}_p \tag{126}$$

As an example, for $N_p = 2$:

$$\vec{\vec{A}}_{-1} = 0\vec{\vec{M}}_0 + \frac{-1/2}{h}\vec{\vec{M}}_1 + \frac{1}{h^2}\vec{\vec{M}}_2, \tag{127}$$

$$\vec{\vec{A}}_0 = 1\vec{\vec{M}}_0 + \frac{0}{h}\vec{\vec{M}}_1 + \frac{-2}{h^2}\vec{\vec{M}}_2, \tag{128}$$

$$\vec{\vec{A}}_1 = 0\vec{\vec{M}}_0 + \frac{1/2}{h}\vec{\vec{M}}_1 + \frac{1}{h^2}\vec{\vec{M}}_2. \tag{129}$$

At the edges of the simulation domain, periodic boundaries are assumed. The spill-over areas from the all-orders model are left intact. At the walls, at $x = \pm a$, the boundary conditions for E_y and E_z are again inserted using Lagrange multipliers.

5.3.3 Truncated Polynomial fit results

The results for the truncated Polynomial model is given in Figure 23. All parameters and profiles are kept the same as for the all-orders model: it is again a JET-like scenario, with 5% hydrogen in deuterium, and 5001 grid points. An 8th order polynomial was used, with a fit window the size of $\zeta = 1.1$.

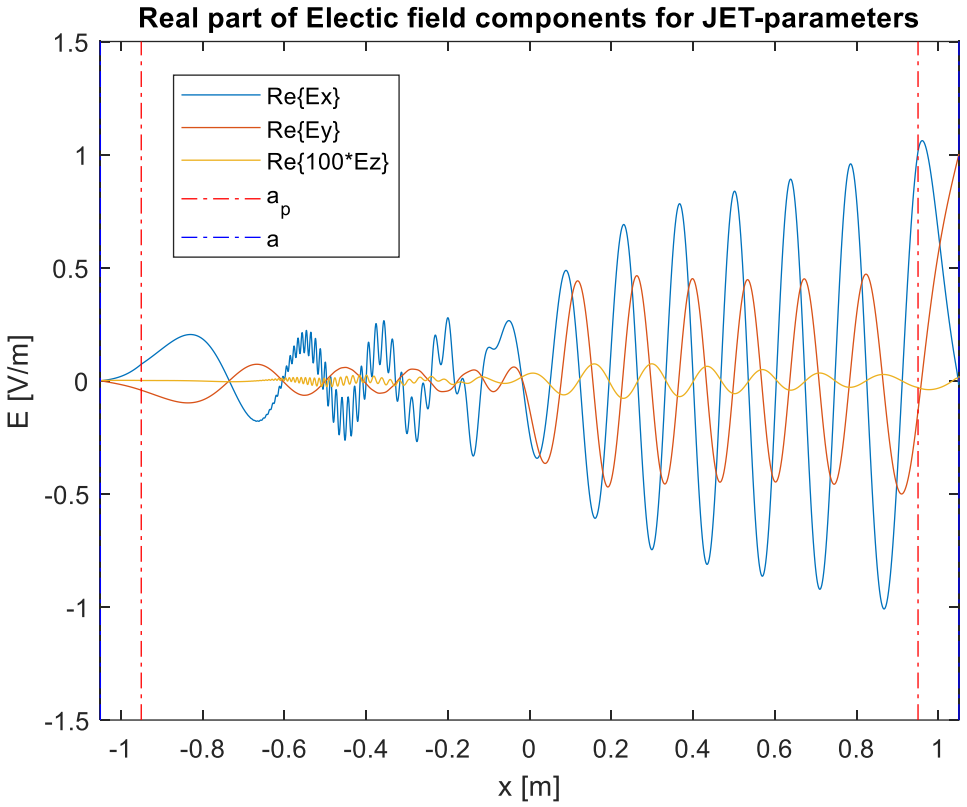

Figure 23: Real part of the electric field components, as calculated by the truncated Taylor Polynomial model. Only the region between $-a$ and a is plotted, bounded by the dashed blue lines, the spill-over regions are cut away. Furthermore, the transition between the core plasma profiles and the edge plasma profiles is indicated with the dashed red lines.

The graph is remarkably similar to all-orders results in Figure 19, to the point where side-by-side comparisons reveal very little differences. Therefore, a quantitative description of the differences is given in the next chapter. However, it should be noted that in contrast to the truncated Taylor model, the IBW follows very similar characteristics to the all-orders model. Visual inspection reveals that the IBW has damped out by the time it reaches $x = -0.69$m, which is very similar to the $x = -0.66$m observed for the all-orders model.

6 Comparison of hot-plasma models

In the previous chapter, three hot-plasma models have been introduced. Visual inspection gives some clues about the differences and similarities, but in order to rigorously compare how the models compare to each other, a metric that quantifies the error is desired. For this purpose, the Root Relative Squared Error (RRSE) is introduced [55]. This metric measures the error in the fields, normalized to the total power in the "signal". When the simulation is denoted as P, and the actual value as A, the $RRSE$ is given by

$$RRSE = \sqrt{\frac{\sum_{n=1}^{N}|P_n - A_n|^2}{\sum_{n=1}^{N}|A_n - \bar{A}|^2}}. \tag{130}$$

In this equation, $\bar{A}$ is the mean of the actual data,

$$\bar{A} = \frac{1}{N}\sum_{n=1}^{N} A_n. \tag{131}$$

The advantage of this metric is that it is provides an error relative to the total signal strength, without magnifying errors around zero-crossings. One disadvantage can be that phase shifts are punished. This is demonstrated in Figure 24: a small phase shift is seen as equally bad as a significant amplitude difference. However it turns out that all models produce very similar fast-wave results, both in phase and magnitude. The RRSE is useful in this case, as it mainly punishes differences in the IBW behavior: when the damping of this wave is not correct, there are often associated shifts present in the IBW as well.

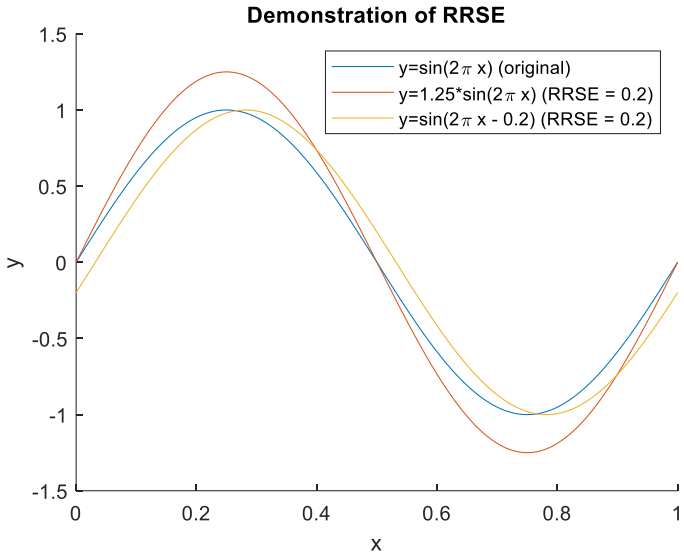

Figure 24: Example showing equal RRSE for a scaled and a shifted version of a sine wave.

Other metrics could be used for additional information. The wave propagation simulations are often run in order to find the answer to some specific question: "where does the injected power end up", "how efficient is the heating strategy", or "which species gains the most energy". This model lacks a description for the kinetic flux, which means that part of the power absorption is not accounted for. As such, these models cannot give a quantitative answer to those questions. Therefore, the choice is made to compare only the results for the electric field through the RRSE. However, section 6.2 gives a qualitative comparison between the absorbed power, based on the Poynting flux alone, between the all-orders model and the truncated Taylor series model.

The subsequent sections of this chapter compare the models with each other, bases on the RRSE metric. The last section concludes with a discussion about the computational time required by the three models.

6.1 Comparison of the electric field generated by the three hot-plasma models

Visual inspection of the electric fields generated by the three hot-plasma models already reveals some differences. In principle there are three electric field components: E_x, E_y and E_z. However, E_z is very weak due to the large conductivity parallel to the magnetic field, and E_y does not contain the mode-converted ion Bernstein wave for the JET-case under consideration. Therefore, only $|E_x|$ is shown in Figure 25. The focus is on the region where the ion Bernstein Wave is present, which is to the left of the fundamental hydrogen ion cyclotron resonance layer near $x = 0$.

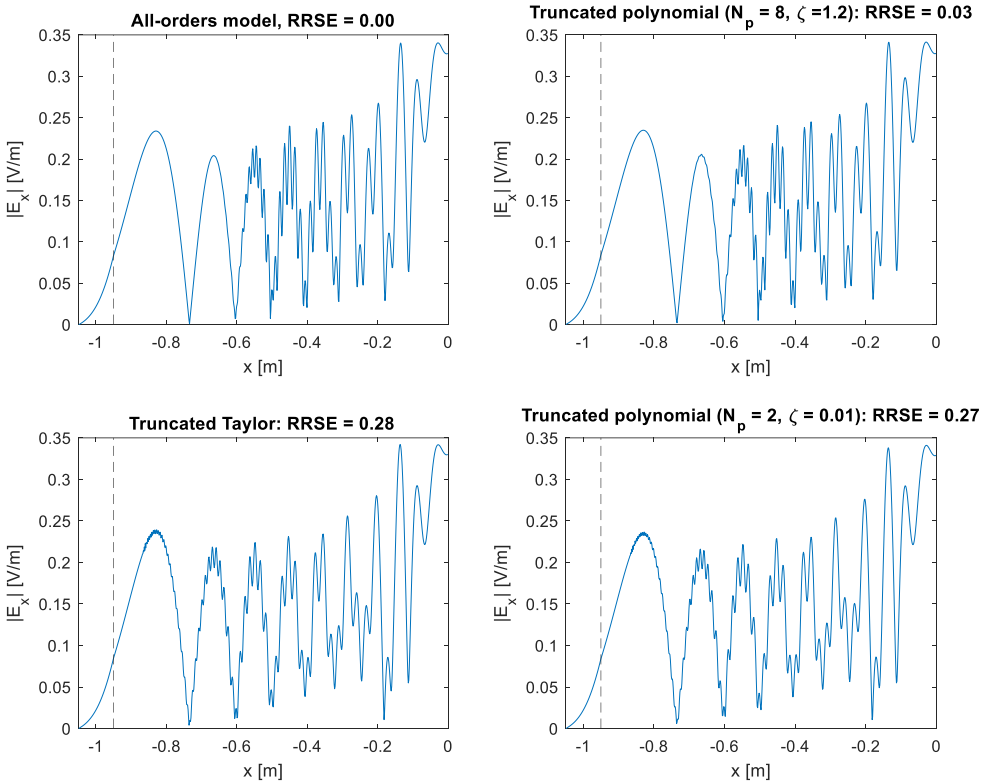

Figure 25: Comparison of the results for $|E_x|$ by each model. The top two graphs show the "all-orders-like behavior" from the all-orders model (top left), and a high-order truncated polynomial fit model (top right). The bottom two graphs show the "truncated FLR-like behavior" from the truncated Taylor series model (bottom left) and a low-order truncated polynomial fit model (bottom right). The black dashed lines represent the separation between core plasma and edge plasma. Note mainly the difference in damping of the short-wavelength wave from right to left.

The main difference between the 4 graphs can be found in the amount of damping of the mode-converted ion-Bernstein wave from right to left; the fast wave behavior is near-identical regardless of the approach chosen. The two very similar graphs at the top reveal that a high-order truncated polynomial fit model is capable of reproducing fields closely resembling those of the all-orders model. Instead results closely resembling the truncated Taylor series model are achieved for a low polynomial order, demonstrating the versatility of the truncated polynomial fit model.

6.2 Absorbed power – All-orders compared to truncated Taylor series

In order to rigorously describe the absorbed power, the model requires a description for the kinetic flux [19]. However, based on the electric fields alone, an estimate on the absorbed power can be made, based on the power carried by the Poynting flux alone [56].

$$P_{abs}(x) = \frac{1}{2} Re\{E^*(x) J_p(x)\} \tag{132}$$

This will lead to slightly unphysical results, with the absorbed power for example oscillating around 0 for the mode-converted ion Bernstein wave. However, it highlights a difference between the two approaches in the damping of the IBW.

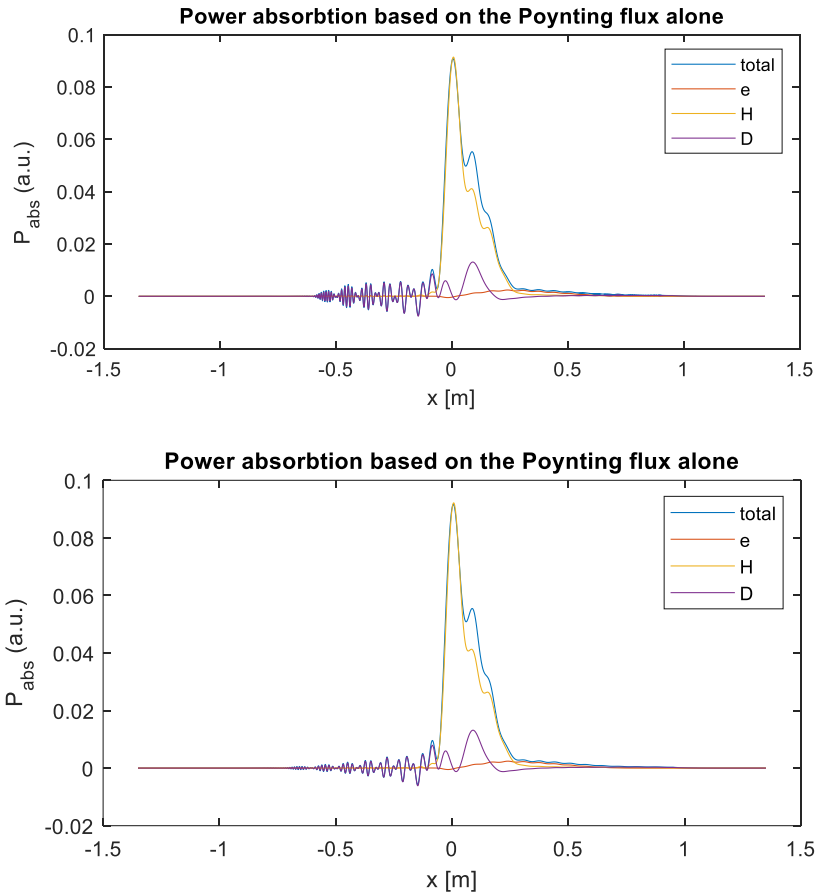

Figure 26: comparison of the absorbed power, based on the Poynting flux, between the all-orders model (top) and the truncated Taylor series model (bottom). Note how in the bottom graph, the power absorption is smeared out over a longer distance.

Note how in the all-orders graph, the power absorption is much more concentrated. This matches the observations in the results for the electric field, where the IBW is damped more quickly. This quicker damping translates to a higher absorption. Depending on the scenario, the differences between the all-orders approach and a truncated FLR approach can be larger than what is shown here [19].

6.3 RRSE as function of N– comparison between models

To show the effect of increasing resolution, a comparison between the three models is made. For the truncated polynomial fit model, the choice is made for a fit window of $\zeta = 1.2$ and a polynomial order of $N_p = 8$. For the JET case under consideration this choice provides all-orders-like results across a large range of resolutions, as demonstrated in Figure 27 and Figure 29.

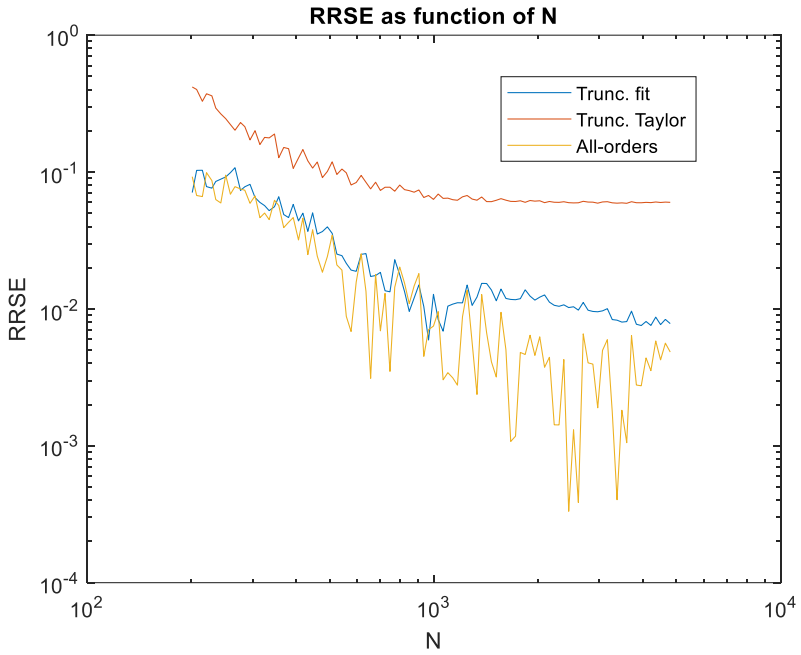

Figure 27: Error as function of the number of grid points, in comparison to a 5000-gridpoints all-orders model. For the truncated polynomial fit model, $N_p = 8$ and $\zeta = 1.2$ is used. Note how the truncated Taylor series model error essentially flatlines after about 1000 grid points, while the truncated polynomial fit model can produce much lower errors.

The truncated Taylor polynomial model does not benefit a lot from increasing the resolution beyond about 1000 grid points. Beyond this point, the fast wave results are correct, but the damping of the Ion Bernstein Wave is not, as demonstrated in Figure 25 . Increasing the resolution does not solve this issue. Meanwhile, the truncated polynomial fit can produce more accurate predictions of the Ion Bernstein Wave damping, which manifests itself in a lower RRSE. Compared to the all-orders model however, it comes at a drastically lower computational cost.

6.4 Comparison of the Computational complexity

For the one-dimensional case under consideration, all models complete within minutes, even the all-orders model. However, as indicated in Table 1 in the introduction, the real benefits are obtained for 2D and 3D situations. To examine how the models compare to each other, the computational time required to obtain the electric field is sampled at 500 different resolutions, and summarized in Figure 28. The computational effort is split in the preparation phase and the actual matrix solve.

For the all-orders model, simulations up to $N = 5000$ have been performed, as this is about the limit for a machine equipped with 16GB RAM. The two finite difference models can be run with much higher resolutions due to the sparsity of the matrices, and as such they have been calculated up to $N = 100\,000$. Linear fits have been calculated through the log-log plot, to determine the N^{α} scaling for both the preparation phase and the matrix solve phase. For the preparation phase, the fit only takes the values for $N > 2000$ into account, as this is where linear behavior starts to be visible in all three models.

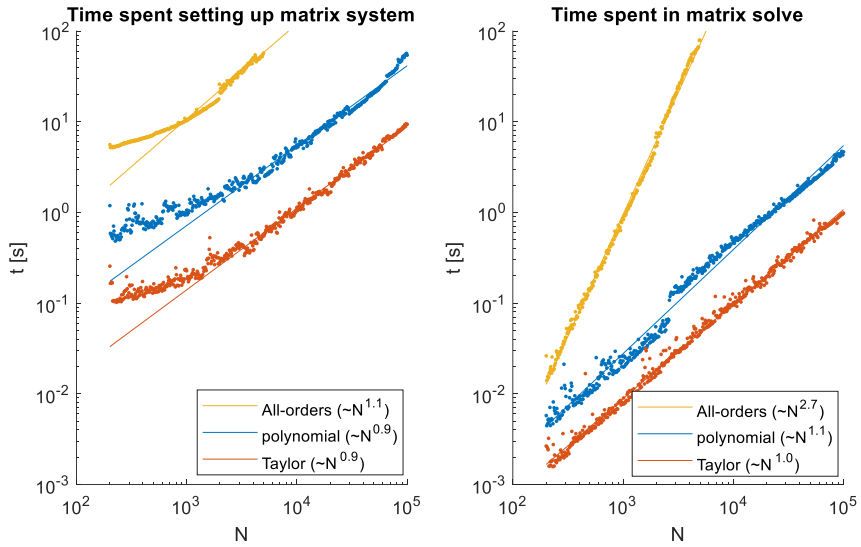

Figure 28: Computational time taken for the electric field calculation by the three different models. The left image shows the time needed for the matrix inversion, while the right image shows the time needed for all other code. The vertical scaling is aligned for easy comparison. The straight lines indicate N^{α} scaling.

For the two finite difference models, the time required to solve the linear system scales approximately with N, while it scales with $N^{2.7}$ for the all-orders model. The preparation phase consumes a significant amount of time for all three models, but scales approximately with N for high resolutions.

In Table 2 these findings are extrapolated to 2D and 3D situations, under the assumption that the same hardware is used and that each direction requires 200 modii. Additionally, the assumption is made that multi-dimensional finite difference schemes influence the computational time by raising the stencil size s to the number of dimensions D: $t \propto s^D$. Without testing in practice, the 2D and 3D values are only speculative, but they should give a rough idea of the relative differences between the models.

	Truncated Taylor	Truncated Polynomial	All-orders
Matrix setup (measured)	$\propto N^{0.9}$	$\propto N^{0.9}$	$\propto N^{1.1}$
Matrix invert (measured.)	$\propto N^{1.0}$	$\propto N^{1.1}$	$\propto N^{2.7}$
Stencil size	3	9	–
Time for $N = 200$ (1D) - measured	0.11 s	0.55 s	5.2 s
Time for $N = 4E4$ (~2D) - approx	$3 \cdot 4$ s = 12 s	$9 \cdot 20$ s = 3 min	5.5 h
Time for $N = 8E6$ (~3D) - approx	$9 \cdot 9$ min = 81 min	$81 \cdot 40$ min = 54 h	920 years

Table 2: Scaling of the computational time for a minimum resolution case (200 modii in each direction) in 1D, and extrapolation to 2D and 3D based on Figure 28. The assumption is made that the 1D computational times are representative for 2D and 3D computational times, and that the same hardware is used for all situations: an HP ZBook studio G4 with a quad-core i7-7700HQ and 16GB RAM.

7 Peculiarities of the truncated polynomial fit model

For the results in the previous chapter, an 8th order polynomial was used with a window size of $\zeta = 1.2$. This choice is not set in stone; there is a lot of freedom in choosing both the polynomial order N_p and the fit window ζ. Both influence the quality of the fit, and the description of the dielectric tensor in k-space. Furthermore, testing revealed that there is one additional, unexpected parameter that influences the results in a non-trivial manner: the number of gridpoints N. Intuition would suggest that higher resolutions yield more accurate simulations, and hence smaller errors. This is not the case for the truncated polynomial fit model: the solution becomes unstable above a certain resolution.

This chapter is dedicated to showing the effect of changing N, ζ and N_p on the RRSE.

7.1 Truncated polynomial fit model – behavior with respect to N and ζ

The results in the previous section showed that the truncated polynomial fit model is able to recreate results closely resembling the all-orders model using appropriate choices for the polynomial order and the fit window. However, further experimentation reveals that besides the polynomial order N_p and the fit window ζ, there is a third parameter that strongly influences the results: the number of grid points N. Figure 29 contains a heatmap of the RRSE as function of both N and ζ for a polynomial of order 8.

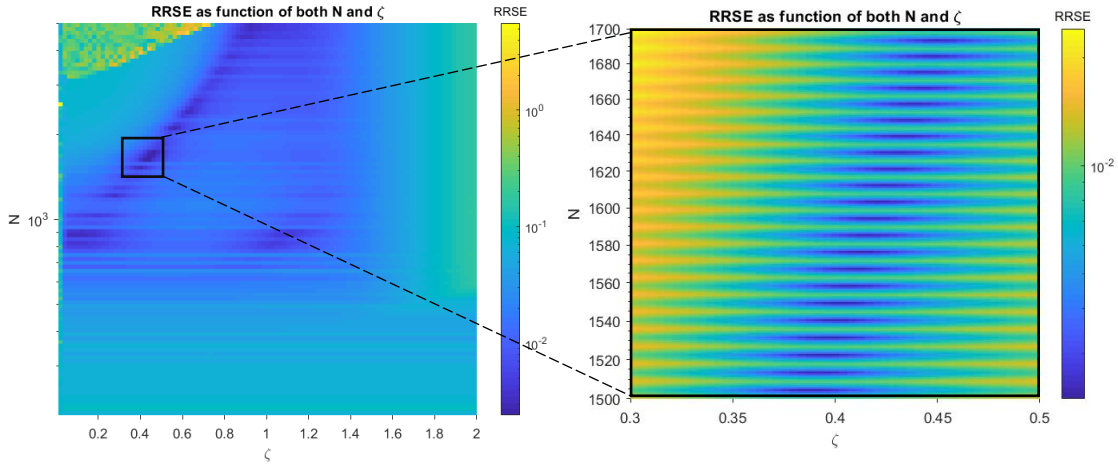

Figure 29: RRSE as function of both the number of grid points N *and the fit window size* ζ*, with an 8th order polynomial. The RRSE is calculated with respect to an all-orders simulation with 5001 grid points. The left picture shows an overview for a large range of* N *and* ζ*, while the right picture highlights a small section in high resolution. Note how in this zoom there are alternating resolutions of high and low RRSE. From the 10.000 samples taken in both figures, a minimum RRSE of 0.22% is found at* $N = 1504$ *and* $\zeta = 0.386$*.*

A clear relation between the resolution, the fit window and the error cannot be derived from this figure. Contrary to expectation, a higher resolution does not automatically lead to a lower error. There is an upper limit to the resolution, above which the error suddenly jumps up. Below this limit, two other interesting cases can be distinguished. Firstly, there are resolutions for which an appropriate selection for the fit window results in a reduced error. Secondly, there are resolutions that perform visibly worse than a slightly higher or slightly lower resolution, regardless of the chosen fit window. This effect is clearest in the zoom in Figure 29, where low-error resolutions alternate with higher-error resolutions.

Tests pointed out that a similar effect can be observed in the truncated Taylor series model, which raises the suspicion that it is not necessarily the fit procedure that produces this behavior. One possible explanation is that the cyclotron resonance layer of one of the species j ends near one of the grid points for some of the resolutions. This would make the argument of the plasma dispersion function ζ_{nj} a small number, see equation (43). In the MATLAB code, the plasma dispersion function is calculated using various approximations for the Dawson function, as the built-in function is very slow especially for $|\zeta_{nj}| < 10$. In this same region the approximations are troublesome as well due to slowly converging sums, potentially resulting in numerical pollution. Tests have been run with the built-in Dawson function and higher precision approximations to no avail.

Other sources of error could be the fact that near the plasma edge, the Larmor radius becomes very small, which means that the fit domain in $k_\perp$ becomes wider than the maximum wave number supported by the resolution $2\pi N/L$. Finally, implementation errors of some kind have not been found, but they cannot be excluded.

These effects combined pose a problem for the application of this model: without a-priori knowledge of the location of the minima in the error, an all-orders model is required to find a suitable combination of the resolution, the fit window and the polynomial order.

7.2 Truncated Polynomial fit model – influence of the polynomial order

As discussed in section 6.1 and section 7.1, the polynomial order N_p is the third non-plasma related parameter that significantly influences the accuracy of the electric field calculation. In order to be able to draw some conclusions on the effect of the polynomial order, the following methodology is used. First the RRSE is calculated for 10.000 (N, ζ) pairs for all polynomial orders between $N_p = 2$ and $N_p = 12$, creating figures like Figure 29a. These results can be found in Appendix F for the interested reader. All (N, ζ) pairs which result in an RRSE>1 are discarded. Of the remaining samples, two figures are extracted: the minimum RRSE and the average RRSE. The results are summarized in Figure 30.

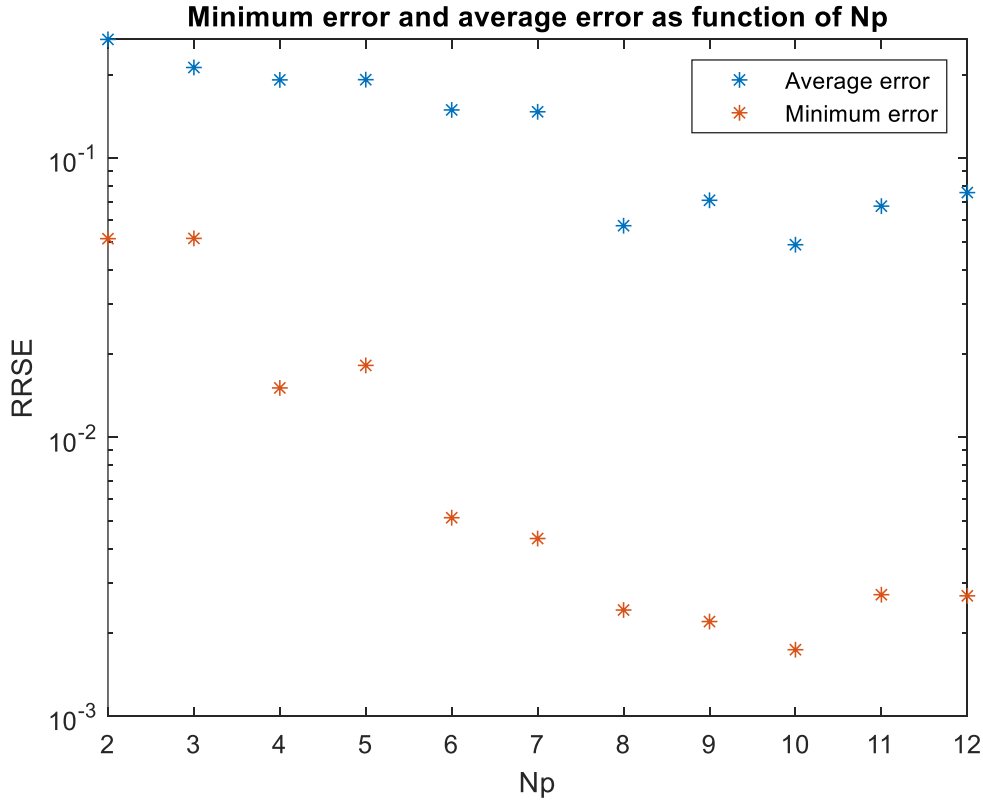

Figure 30: Average RRSE and minimum RRSE for the truncated polynomial fit model, under different polynomial orders. In general, both the average error and the minimum error tend to go down with increasing polynomial order, with the most significant gains up to $N_p = 8$.

In general, there is a clear downward trend with increasing polynomial order for both the minimum RRSE and the average RRSE. This trend continues until at least an 10th order polynomial for the JET case under consideration. After this point the error seems to increase again, although no datapoints were generated beyond $N_p = 12$.

Furthermore, it seems to be beneficial to use polynomials of even order. This could be attributed to usage of the central difference scheme, which uses only symmetric stencils of odd sizes. Therefore, an odd polynomial, say $N_p = 7$ uses the same stencil size as the next even polynomial, $N_p = 8$. For nearly no added computational resources, a better fit can be created, resulting in a lower error.

8 Discussion

The truncated polynomial fit model has a clear advantage over the all-orders model in terms of computational complexity. Like the truncated Taylor series model, a N scaling is obtained for the matrix inverse, instead of an $N^{2.7}$ scaling for the all-orders model. This translates to significant savings on the computational resources required to run the simulations, especially for hypothetical 2D and 3D cases. The advantage over the truncated Taylor series model is that it can correctly predict the damping of the Ion Bernstein Wave. This was demonstrated in Figure 25 and Figure 27, both by visual inspection and through the usage of the RRSE metric.

One issue that affects this model, is the fact that there are three non-plasma-related parameters that strongly influence the model performance: the number of grid points N, the window ζ over which the polynomial fit is performed, and the polynomial order N_p. By carefully choosing a suitable combination of these parameters, all-orders results can be achieved with a differential-only approach. The problem is that it is difficult to assess a priori which combination of parameters is most suitable. Finding the right combination as of now requires comparison with an all-orders model, a limitation that dramatically reduces the usability of this model.

An observation from working with this code, is that the system matrices suffer from very high condition numbers, typically of the order $10^{17} - 10^{19}$. The source of this poor conditioning has not been identified yet. As the problem gets worse for high polynomial orders, it might be related to incompatibility between the finite difference scheme and high-order derivatives. Other numerical schemes like finite-elements have not been investigated, perhaps a solution lies in this direction.

Some preliminary testing with quad-precision has been done using the Advanpix Multiprecision Computing Toolbox [57] for MATLAB. Mixed results were obtained; even though quad-precision seemed to bring improvements, it was not a silver bullet that took away the problem entirely. Moreover, the simulations got significantly slower due to MATLAB's poor optimization for multi-precision computing, combined with the fact that current computer hardware has no native support for quadruple precision operations.

There are some aspects of the model that could be improved on by future work (chapter 9). Transition from the finite difference scheme to a finite element scheme would be desired, as it is better suited for non-uniform meshes. It is for example not unthinkable that a high resolution is assigned specifically to the region where short-wavelength waves, like the IBW, are present. This would allow for an easier extension to two dimensions or three dimensions. Additionally, the model currently does not take the DC plasma current into account, which produces an additional poloidal magnetic field. In this case, the assumption that $k_z = k_\parallel$ is not valid anymore. Adding this contribution requires for example a coordinate transformation [23]. A proper description of the kinetic flux would allow for more accurate predictions of the power absorption and current drive.

However, even though the model has its limitations, perhaps the most important conclusion is the fact that this work proves that full-domain solutions like AORSA's spectral collocation method are not always required for accurate predictions of the wave behavior in fusion plasmas. An approach that has a finite width both in the spatial domain and the spectral domain can produce very accurate results, for only a fraction of the computational effort. This opens the door for other approaches that take both the behavior in the spatial domain and the spectral domain into account, like the spectral element method [58] and wavelets [59].

9 Future work

In order to increase the applicability of the model, research should be conducted to determine the exact relation between resolution, fit window and polynomial order. If the results from Figure 29 can be explained by theory, or if a source for the difficult relation can be found, a selection for N_p, ζ and N can be made without needing to compare to an all-orders model. Otherwise, the benefit from the fast model is diminished, as the slow all-orders model always needs to be used to determine whether the results are correct. It would be interesting to test whether the problem persists with other numerical schemes, like the finite element scheme. This scheme has the added advantage that it is easier to deal with non-uniform grids, or higher-dimensional meshes. In this same spirit, it would be interesting to investigate other semi-localized approaches like wavelets and the spectral element method.

A very interesting follow-up question would be to investigate whether this approach can be ported to a 2D or 3D case. Instead of fitting a polynomial of through $\vec{\epsilon}(k_x)|_{x=x_t}$, this would translate to fitting curved planes through $\vec{\epsilon}(k_x, k_y)|_{x,y=x_t,y_t}$ in a 2D simulation. The fitting procedure itself is not very difficult, as demonstrated for a simple case below ($N_{px} = 4, N_{py} = 4$). The question is whether this can be used to produce accurate 2D or even 3D simulations.

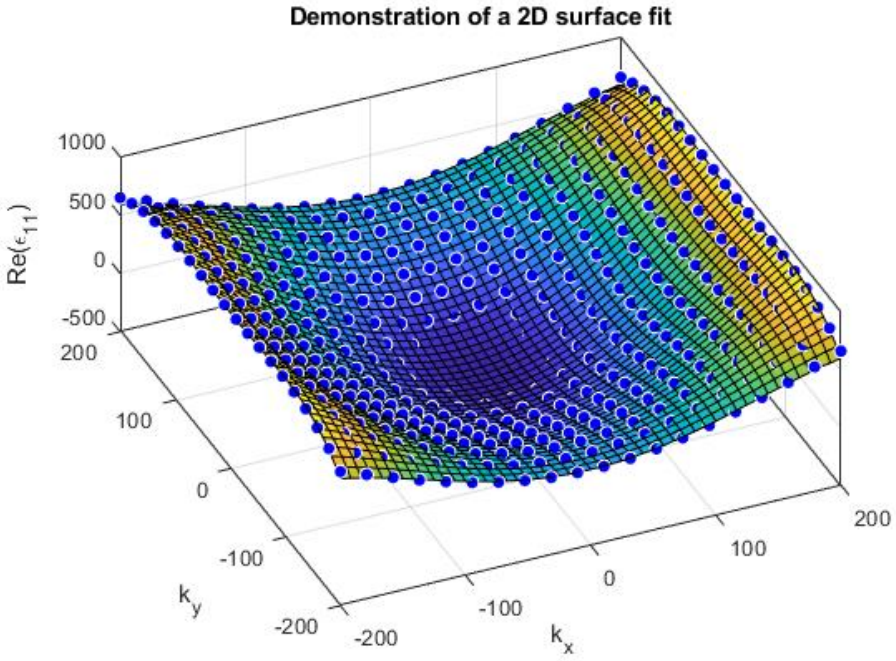

Figure 31: Demonstration of a 2D surface fit of $Re\{\epsilon_{11}\}$, for a 4th order polynomial in both X and Y direction.

Even if this approach does not scale to a two- or three-dimensional case, a hybrid approach can be considered. This approach would use the truncated polynomial fit approach in the radial direction, and the spectral collocation method in poloidal and toroidal direction. The motivation for this choice is that the radial direction shows the strongest aperiodic behavior. Sources of the aperiodicity are evanescent behavior in the plasma edge; one-sided excitation, and presence of mode-conversion; reflection; and absorption layers. The poloidal and toroidal directions however are a lot more periodic, and the expectation is that these directions are better suited for a spectral method. The hybrid case will have reduced computational costs compared to using the spectral method in all directions.

Finally, in order to make conclusions about the power deposition location and power fluxes, a proper description for the kinetic flux needs to be included in the model. The TOMCAT and TOMCAT-U papers by Dirk van Eester discuss how to correctly account for the kinetic flux but retrofitting this to the existing model is not feasible. Starting from scratch with the additional equations is the preferable route.

10 Conclusion

Simulating the behavior of electromagnetic waves in hot, magnetized fusion plasmas is a challenging numerical problem. Especially correctly predicting the behavior of the Ion Bernstein Wave has been proven to be a non-trivial task. In order to get accurate predictions for this wave, up until now all-orders hot plasma models were used, that don't make assumptions on the smallness of the Larmor radius compared to the wavelength under investigation. The typical approach in this case is a spectral collocation method, which results in computationally demanding simulations, especially for 2D and 3D cases.

In order to investigate whether there is a method that reproduces accurate predictions for the IBW damping, without extreme computational expenses, four one-dimensional models have been created. First, a cold plasma model is discussed. This model does not contain the IBW, but it serves as a starting point for the more complex hot plasma models. Three hot plasma models were subsequently developed.

The truncated Taylor series model is the fastest model of the three. The assumption is made that the spectrum of the dielectric tensor is very narrow, such that a Taylor approximation of the dielectric tensor is allowed around $k_x = 0$. Using a very useful property of the Fourier transform and some mathematical manipulations, a sparse and banded matrix is obtained. The assumption on the narrowness of the spectrum results in a correct prediction of waves with a low wavenumber, such as the fast wave. On the other hand, high-wavenumber waves like the Ion Bernstein Wave are not predicted correctly, as is visible in Figure 25. This is a general conclusion for truncated Finite Larmor Radius codes; the damping of the Ion Bernstein Wave is underestimated. Quantitative comparisons of the absorbed power confirmed these findings.

The slowest but most accurate model is the all-orders model, a model that makes no finite Larmor Radius expansions. The AORSA code from Fred Jaeger is used as a starting point, and it is adapted for 1-dimensional application. AORSA is often considered to be a benchmark for other numerical schemes, and the same vision is adopted in this work. This model is the slowest of the three hot plasma models. A very strong scaling of the computational effort with the resolution is found: $t \propto N^{2.7}$. Therefore, 2D simulations with the AORSA code are often run on computing clusters with 100's or 1000's of CPU cores, let alone the resources required for a high-resolution 3D simulation.

The focus of this work has been on finding a mid-way between both approaches, combining the speed of the truncated Taylor series model with the accuracy of the all-orders model. In order to accomplish this, the truncated Taylor series model has been adjusted, forming the truncated polynomial fit model. Instead of forming a Taylor series around $k_x = 0$, a polynomial of powers of k_x is fitted through $\epsilon_{x=x_t}(k_x)$. It is shown that this approach is capable of producing results that closely resemble the all-orders results, at a computational cost that scales with N, instead of $N^{2.7}$. Especially for 2D and 3D situations, the performance gain could be significant.

Some issues with this approach have been identified. Mainly the complex relation between the resolution, the fit window and the polynomial order is currently not understood. As such, it is not possible to estimate a priori which combination of these parameters yields a low-error result. Some possible ways forward have been discussed. This includes investigating other numerical schemes that might be more robust.

To conclude, this approach demonstrates that it is possible to achieve all-orders results without a full-domain approach like the all-orders model. It opens the door for other semi-localized wave representations, like wavelets or the spectral element method. Examining these options and extending the idea to 2D would be interesting follow-up research.

11 References

[1] G. Maggio and G. Cacciola, "When will oil, natural gas, and coal peak?," *fuel*, vol. 98, pp. 111-123, March 2012.

[2] IPCC, "Climate Change 2014: Synthesis Report. Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change [Core Writing Team, R.K. Pachauri and L.A. Meyer (eds.)]," IPCC, Geneva,Switzerlad, 2014.

[3] OACD/IEA, "CO2 Emissions from Fuel Combustion 2018 Highlights," IEA, 2018.

[4] International Energy Agency (IEA), "Key world energy statistics," IEA, 2018.

[5] J. Ongena and Y. Ogewa, "Nuclear fusion: Status report and future prospects," *Energy Policy*, vol. 96, no. 770-778, 2016.

[6] E. G. Adelberger et al, "Solar fusion cross sections. II. The pp chain and CNO cycles," *Review of Modern Physics*, vol. I, no. 83, pp. 195-245, 2011.

[7] E. Morse, "4.5 Energy Balance and Lawson Criterion," in *Nuclear Fusion*, Switzerland, Springer, 2019, pp. 97-100.

[8] Dstrozzi, "File:Fusion tripleprod.svg," 26 November 2010. [Online]. Available: https://en.wikipedia.org/wiki/File:Fusion_tripleprod.svg. [Accessed 22 July 2019].

[9] J. R. De Laeter, J. K. Böhlke, P. De Bièvre, H. Hidaka, H. S. Peiser, K. J. R. Rosman and P. D. P. Taylor, "Atomic weights of the elements: review 2000," *Pure Applied Chemistry*, vol. 75, no. 6, pp. 683-800, 2003.

[10] L. L. Lucas and M. P. Unterweger, "Comprehensive Review and Critical Evaluation of the Half-Life of Tritium," *Journal of Research of NIST*, vol. 105, no. 4, pp. 541-549, 2000.

[11] M. Nishikawa, "Tritium Breeding Ratio required to keep Tritium balance in a D-T Fusion Reactor," *Fusion Science and Technology*, vol. 60, no. 3, pp. 1071-1076, 2011.

[12] P. Grammelis, N. Margaritis and E. Karampinis, "2 - Solid fuel types for energy generation: Coal and fossil carbon-derivative solid fuels," in *Fuel Flexible Energy Generation*, Woodhead Publishing, 2016, pp. 29-58.

[13] P. W. Gruber, P. A. Medina and G. A. Keoleian et al, "Global Lithium availablility," *Journal of Industrial Ecology*, 2011.

[14] S. Rosanvallon, B. C. Na, M. Benchikhoune, J. Elbez Uzan, O. Gastaldi, N. Taylor and L. Rodriguez, "ITER waste management," *Fusion Engineering and Design*, vol. 85, no. 10-12, pp. 1788-1791, 2010.

[15] ITER, "ITER Goals," ITER, [Online]. Available: https://www.iter.org/sci/Goals. [Accessed 23 July 2019].

[16] Conleth Brady / IAEA, "ITER Exhibit (01810402) - ITER exhibits at the International Fusion Energy Days 2013. Monaco. 2 December 2013," 2 December 2013. [Online]. Available: https://www.flickr.com/photos/iaea_imagebank/12219071813/. [Accessed 23 July 2019].

[17] J. Freidberg, "Power balance in a fusion reactor," in *Plasma Physics and Fusion Energy*, New York, Cambridge University Press, 2007, pp. 60-82.

[18] E. A. Aziziov, "Tokamaks: from A. D. Sakharov to the present (the 60-year history of tokamaks)," *Physics-Uspeki*, vol. 55, no. 2, pp. 190-203, 2012.

[19] D. van Eester and E. Lerche, "A 1D model for describing ion cyclotron resonance heating at arbitrary cyclotron harmonics in tokamak plasmas," *Plasma Physics and Controlled Fusion*, vol. 55, no. 5, 2013.

[20] J. C. Wright, J. Lee, E. Valeo, P. Bonoli, C. K. Phillips, E. F. Jaeger and R. W. Harvey, "Challenges in Self-Consistent Full-Wave Simulations of Lower Hybrid Waves," *IEEE Transactions on Plasma Science*, vol. 38, no. 9, pp. 2136-2143, 2010.

[21] J. C. Wright, B. T. Bonoli, F. M. M. Brambilla, E. D'Azevedo, D. B. Batchelor, E. F. Jaeger, L. A. Berry, C. K. Phillips and A. Pletzer, "Full Wave Simulations of Fast Wave Mode Conversion and Lower Hybrid Wave Propagation in Tokamaks," *Physics of Plasmas*, vol. 11, no. 5, 2004.

[22] D. L. Green and L. A. Berry, "Iterative addition of parallel temperature effects to finite-difference simulation of radio-frequency wave propagation in plasmas," *Computer Physics Communications*, vol. 185, pp. 736-743, 2014.

[23] E. F. Jaeger, L. A. Berry, E. D'Azevedo, D. Batchelor and M. D. Carter, "All-orders spectral calculation of radio-frequency heating in two-dimentional toroidal plasmas," *Physics of Plasmas,* vol. 8, no. 5, pp. 1573-15783, 2001.

[24] E. D'Azevedo and J. C. Hill, "Parallel LU Factorization on GPU Cluster," *Precedia Computer Science,* vol. 9, pp. 67-75, 2012.

[25] Intel, "Intel Optane Memory," Intel, [Online]. Available: https://www.intel.com/content/www/us/en/architecture-and-technology/optane-memory.html. [Accessed 25 july 2019].

[26] R. J. Dumont, "Variational approach to radiofrequency waves in magnetic fusion devices," *Nuclear Fusion,* vol. 49, p. 075033, 2009.

[27] D. van Eester and R. Koch, "A variational principle for studying fast-wave mode conversion," *Plasma Physics and Controlled Fusion,* vol. 40, pp. 1949-1975, 1998.

[28] J. Meichsner, M. Schmidt, R. Schneider and H.-E. Wagner, "Thermal and Nonthermal plasmas," in *Nonthermal Plasma Chemistry and Physics*, Miami, Taylor & Francis Group, 2013, pp. 33-35.

[29] D. G. Swanson, "Waves in a cold uniform plasma," in *Plasma Waves - Second Edition*, Bristol and Philidelphia, Institute of Physics Publishing, 2003, pp. 35-93.

[30] D. G. Swanson, "Properties of plasmas," in *Plasma Waves, 2nd Edition*, Bristol and Philadelphia, Institute of Physics Publishing, 2003, pp. 15-17.

[31] D. B. Batchelor et al, "Electromagnetic mode conversion: understanding waves that suddenly change their nature," *Journal of Physics: Conference Series,* vol. 16, pp. 35-39, 2005.

[32] J. A. Bittencourt, "Magnetohydrodynamic waves," in *Fundamentals of Plasma Physics*, New York, Springer, 2004, pp. 375-397.

[33] D. van Eester, F. Louche and R. Koch, "Re-evaluation of ITER ion cyclotron operating scenarios," *Nuclear Fusion,* no. 42, pp. 310-328, 2002.

[34] C. Castaldo and A. Cardinali, "Tritium Minority Heating With Mode Conversion Of Fast Waves," *AIP Conference Proceedings,* no. 1069, pp. 265-270, 2008.

[35] W. H. Hayt Jr. and J. A. Buck, "Time-Varying Fields and Maxwell's Equations," in *Engineering Electromagnetics - 6th ed.*, Mc Graw Hill, 2001, pp. 322-347.

[36] G. Cattanei, "Resonant Faraday shield ICRH antenna," *Nuclear Fusion,* vol. 42, pp. 541-546, 2002.

[37] T. H. Stix, "Waves in a Cold Uniform Plasma," in *Waves in Plasmas*, New York, American Institute of Physics, 1992, pp. 25-40.

[38] D. G. Swanson, "Partionally ionized plasmas and collisions," in *Plasma Waves - 2nd edition*, Bristol and Philadelphia, Institute of Physics Publishing, 2003, pp. 117-121.

[39] D. G. Swanson, "Waves in a magnetized hot plasma," in *Plasma Waves*, Bristol and Philedelphia, Institute of Physics Publishing, 2003, pp. 176-196.

[40] T. H. Stix, "Solution of the Vlasov Equation," in *Waves in Plasmas*, New York, AIP-Press, 1992, pp. 247-250.

[41] D. G. Swanson, "The Vlasov equations," in *Waves in plasma*, New York, Academic Press, 2003, pp. 145-146.

[42] J. A. Bittencourt, "Waves in hot isotropic plasmas," in *Fundamentals of Plasma Physics*, New York, Springer, 2004, pp. 483-514.

[43] D. van Eester and E. Lerche, "Kinetic Theory of Plasma Waves," *Transactions of Fusion Science and Technology,* vol. 61, no. 2T, pp. 331-345, 2012.

[44] D. G. Swanson, "The evolution of the distribution function," in *Plasma Waves*, Bristol, Institute of Physics Publishing, 2003, pp. 177-180.

[45] D. G. Swanson, "Appendix B: Special functions in plasma physics," in *Plasma Waves, 2nd Edition*, Bristol and Philadelphia, Institute of Physics Publishing, 2003, pp. 415-427.

[46] M. Brambilla, "The constitutive relation," in *Kinetic Theory of Plasma Waves*, Clarindon Press, 1998, pp. 17-23.

[47] Wikipedia, "Torus," Wikipedia, 3 July 2019. [Online]. Available: https://en.wikipedia.org/wiki/Torus. [Accessed 11 July 2019].

[48] C. Sborchia, A. B. Olivia and T. Boutkoul et al., "The ITER magnet systems: progress on construction," *Nuclear Fusion,* vol. 54, 2014.

[49] W. Dahmen and A. Kunoth, "Appending Boundary Conditions by Lagrange Multipliers: Analysis of the LBB Condition," *Numerische Mathematik,* vol. 88, no. 1, 1999.

[50] D. van Eester, "Notes on solving wave equation using Fourier method [unpublished]," Laboratorium voor Plasmafysica - Laboratoire de Physique des Plasmas, 2018.

[51] C. Grossmann and H.-G. Roos, " Finite Difference Methods," in *Numerical Treatment of Partial Differential Equations*, Berlin, Heidelberg, New York, Springer , 2007, pp. 23-30.

[52] Wikipedia, "Finite difference coefficient," Wikipedia, 18 07 2019. [Online]. Available: https://en.wikipedia.org/wiki/Finite_difference_coefficient. [Accessed 18 08 2019].

[53] MathWorks, "fft," Mathworks, [Online]. Available: https://nl.mathworks.com/help/matlab/ref/fft.html. [Accessed 25 6 2019].

[54] Wikipedia, "Tables of important Fourier transforms," Wikipedia, 2019. [Online]. Available: https://en.wikipedia.org/wiki/Fourier_transform#Tables_of_important_Fourier_transforms. [Accessed 16 8 2019].

[55] I. H. Witten and E. Frank, "Evaluating numeric prediction," in *Data Mining, Practical Machine Learning Tools and Techniques*, Elsevier, 2005, pp. 176-179.

[56] D. N. Smithe, "Local full-wave energy and quasilinear analysis in nonuniform plasmas," *Plasma Physics and Controlled Fusion,* vol. 31, no. 7, pp. 1105-1126, 1989.

[57] Advanpix, "Multiprecision Computing Toolbox for MATLAB," Advanpix, 30 June 2019. [Online]. Available: https://www.advanpix.com/. [Accessed 29 July 2019].

[58] D. A. Kopriva, "Spectral Element Methods," in *Implementing Spectral Methods for Partial Differential Equations*, Springer Science + Business Media B.V, 2009, pp. 293-347.

[59] P. Vallejos, T. Johnson and T. Hellsten, "Modeling RF waves in spatially dispersive inhomogeneous plasma using an iterative wavelet spectral method," *European Physical Journal Conferences,* no. 157, 2017.

[60] Math Insight, "Introduction to Taylor's theorem for multivariable functions," [Online]. Available: https://mathinsight.org/taylors_theorem_multivariable_introduction. [Accessed 14 12 2018].

[61] J. D. Lawson, "Some Criteria for a Power Producing Thermonuclear Reactor," *Proceedings of the Physical Society. Section B,* vol. 70, no. 1, pp. 6-10, 1957.

[62] P. Li, "Simulation of Electromagnetic Waves in the Magnitized Cold Plasma by a DGFETD Method," *IEEE Antennas and Wireless Propagation Letters,* vol. 12, no. 1, pp. 1244-1247, 2013.

Appendix A: Generalized dielectric tensor

The cold plasma dielectric tensor that is commonly found in literature is generally derived for magnetic fields aligned in the $\hat{z}$-direction. In this section a general dielectric tensor is derived, following a similar approach to Swanson, waves in plasma , starting from the linearized equation of motion for a single particle:

$$-i\omega m_j \vec{v}_{1j} = q_j\left(\vec{E}_1 + \vec{v}_{1j} \times \vec{B}_0\right) \tag{133}$$

For an arbitrary $\vec{B}_0$, this translates to the system of equations:

$$\begin{bmatrix} -i\omega m_j & -B_{0,z}q_j & B_{0,y}q_j \\ B_{0,z}q_j & -i\omega m_j & -B_{0,x}q_j \\ -B_{0,y}q_j & B_{0,x} & -i\omega m_j \end{bmatrix} \begin{bmatrix} v_x \\ v_y \\ v_z \end{bmatrix} = \begin{bmatrix} q_j E_x \\ q_j E_y \\ q_j E_z \end{bmatrix} \tag{134}$$

This system can be solved for $\vec{v}_{1j}$ using standard linear algebra. To clean up the equations, the cyclotron frequency is written in vector form as $\vec{\omega}_{cj}$, and its $\hat{x}$-, $\hat{y}$- and $\hat{z}$- contributions ω_{cjx}, ω_{cjy} and ω_{cjz} are introduced.

$$\vec{\omega}_{cj} = \begin{bmatrix} \omega_{cjx} \\ \omega_{cjy} \\ \omega_{cjz} \end{bmatrix} = \frac{|q_j|}{m_j} \begin{bmatrix} |B_{0,x}| \\ |B_{0,y}| \\ |B_{0,z}| \end{bmatrix} \tag{135}$$

$$\omega_{cj}^2 = \left|\vec{\omega}_{cj}\right|^2 = \omega_{cjx}^2 + \omega_{cjy}^2 + \omega_{cjz}^2 \tag{136}$$

This gives for $\vec{v}_{1j}$, temporarily dropping the subscript j for compactness:

$$\begin{bmatrix} v_x \\ v_y \\ v_z \end{bmatrix} = \frac{iq}{m\omega(\omega_c^2 - \omega^2)} \begin{bmatrix} \omega_{cx}^2 - \omega^2 & \omega_{cx}\omega_{cy} + i\omega_{cz}\omega & \omega_{cx}\omega_{cz} + i\omega_{cy}\omega \\ \omega_{cy}\omega_{cx} + i\omega_{cz}\omega & \omega_{cy}^2 - \omega^2 & \omega_{cy}\omega_{cz} - i\omega_{cx}\omega \\ \omega_{cz}\omega_{cx} - i\omega_{cy}\omega & \omega_{cz}\omega_{cy} + i\omega_{cx}\omega & \omega_{cz}^2 - \omega^2 \end{bmatrix} \begin{bmatrix} E_x \\ E_y \\ E_z \end{bmatrix} \tag{137}$$

Or, in a more compact notation, and re-introducing the subscript j:

$$\begin{bmatrix} v_{1jx} \\ v_{1jy} \\ v_{1jz} \end{bmatrix} = \frac{iq_j}{m_j\omega\left(\omega_{cj}^2 - \omega^2\right)} \left(\vec{\omega}_{cj}\vec{\omega}_{cj}^T + i\omega \begin{bmatrix} i\omega & -\omega_{cjz} & \omega_{cjy} \\ \omega_{cjz} & i\omega & -\omega_{cjx} \\ -\omega_{cjy} & +\omega_{cjx} & i\omega \end{bmatrix} \right) \begin{bmatrix} E_x \\ E_y \\ E_z \end{bmatrix} \tag{138}$$

From this expression for $\vec{v}_{1j}$, $\vec{\vec{\sigma}}$ can be calculated using:

$$\vec{J} = \sum_j n_j q_j \vec{v}_j = \vec{\vec{\sigma}} \cdot \vec{E} \tag{139}$$

The result is:

$$\vec{\vec{\sigma}} = \sum_j \frac{i\epsilon_0\omega_p^2}{\omega(\omega_c^2 - \omega^2)} \left(\vec{\omega}_{cj}\vec{\omega}_{cj}^T + i\omega \begin{bmatrix} i\omega & -\omega_{cjz} & \omega_{cjy} \\ \omega_{cjz} & i\omega & -\omega_{cjx} \\ -\omega_{cjy} & +\omega_{cjx} & i\omega \end{bmatrix} \right) \tag{140}$$

Using (18), this gives the dielectric tensor $\vec{\vec{\epsilon}}$ for an arbitrarily directed $\vec{B}_0$:

$$\vec{\vec{\epsilon}} = \vec{\vec{I}} - \sum_j \frac{\omega_{pj}^2}{\omega^2\left(\omega_{cj}^2 - \omega^2\right)} \left(\vec{\omega}_{cj}\vec{\omega}_{cj}^T + i\omega \begin{bmatrix} i\omega & -\omega_{cjz} & \omega_{cjy} \\ \omega_{cjz} & i\omega & -\omega_{cjx} \\ -\omega_{cjy} & +\omega_{cjx} & i\omega \end{bmatrix} \right) \tag{141}$$

Appendix B: JET parameters

Parameter	Value	Description
N	2501	number of grid points in k-space
$N_{harm,max}$	4	(symmetric) number of terms in the dielectric tensor
N_k	50	number of k-space sample points.
N_p	8	polynomial order
ζ	1.2	k_perp*rho_L number for throughout domain
R_0	2.97 m	major radius
a_s	1.35 m	Total size of simulation domain
x_{ant}	1.05 m	antenna position relative to machine axis
a_p	0.95 m	plasma minor radius
ν/ω	1e-3	collision frequency relative to omega
f	51e6 Hz	Antenna frequency
n_{tor}	27	toroidal mode number (kz = ntor/r)
k_{y0}	0	Wavenumber in y-direction
curve_k_parr	true	make k_parr dependent on r as ntor/r
aperature	true	use aperature-type excitation (1) or current sheet (0)
B_0	3.45 T	magnetic field on axis
n_0	7e19 m^{-3}	edge density, if a parabolic profile is used
$n_{0,edge}$	2e19 m^{-3}	edge density, if a parabolic profile is used
α_n	1	density profile parameter if a parabolic profile is used
λ_n	0.05	tune exponential decay in edge (smaller = faster decay)
T_0	5e3 eV	core temp (eV)
$T_{0,edge}$	1e2 eV	edge temperature (eV)
α_T	1.5	profile shape parameter for temperature
λ_T	0.05	tune exponential decay in edge (smaller = faster decay)
$N_{species}$	2	number of ion species
A_{ion}	[1 2]	mass number of ions
Z_{ion}	[1 1]	charge of ions
$conc_{ion}$	[0.05 0.95]	concentration relative to ne

The density and temperature profiles of each ion species can be set separately from the electrons, but in this case the profiles are assumed to be equal for both the ions and the electrons.

Appendix C: Curved magnetic field lines

To account for the background magnetic field curvature, a plane wave solution is introduced in a cylindrical coordinate system. The magnetic field is aligned in the $\hat{\phi}$-direction, while the variation is in $\hat{\rho}$-direction.

$$\vec{E}(\vec{r}) = \vec{E}(\rho)e^{i(k_\phi \phi + k_z z)} \tag{142}$$

Expressing the double curl in the Helmholtz equation as a matrix product yields:

$$\begin{aligned} &-\nabla \times \nabla \times \vec{E} \\ &= \begin{bmatrix} \frac{1}{r^2}\frac{\partial^2}{\partial \phi^2} + \frac{\partial^2}{\partial z^2} & -\frac{1}{r^2}\frac{\partial}{\partial \phi} - \frac{1}{r}\frac{\partial^2}{\partial \phi \partial r} & -\frac{\partial^2}{\partial z \partial r} \\ \frac{1}{r^2}\frac{\partial}{\partial \phi} - \frac{1}{r}\frac{\partial^2}{\partial \phi \partial r} & \frac{\partial^2}{\partial z^2} + \frac{\partial^2}{\partial r^2} + \frac{1}{r}\frac{\partial}{\partial r} - \frac{1}{r^2} & -\frac{1}{r}\frac{\partial^2}{\partial \phi \partial z} \\ -\frac{1}{r}\frac{\partial}{\partial z} - \frac{\partial^2}{\partial r \partial z} & -\frac{1}{r}\frac{\partial^2}{\partial \phi \partial z} & \frac{\partial^2}{\partial r^2} + \frac{1}{r}\frac{\partial}{\partial r} + \frac{1}{r^2}\frac{\partial^2}{\partial \phi^2} \end{bmatrix} \begin{bmatrix} E_r \\ E_\phi \\ E_z \end{bmatrix} \end{aligned} \tag{143}$$

Comparing this tensor with the Cartesian expression is a little challenging, as the $\hat{\phi}$ direction in the cylindrical system corresponds to the $\hat{z}$ direction of the Cartesian system. Hence, the components are reshuffled such that they look more like $\overrightarrow{\overrightarrow{D}}$ from (68).

$$\begin{bmatrix} \left(-\nabla \times \nabla \times \vec{E}\right)_\rho \\ \left(-\nabla \times \nabla \times \vec{E}\right)_z \\ \left(-\nabla \times \nabla \times \vec{E}\right)_\phi \end{bmatrix} = \overrightarrow{\overrightarrow{D}} \cdot \begin{bmatrix} \vec{E}_\rho \\ \vec{E}_z \\ \vec{E}_\phi \end{bmatrix} \tag{144}$$

This gives the expression for $\overrightarrow{\overrightarrow{D}}$, without plane wave expansion:

$$\overrightarrow{\overrightarrow{D}} = \begin{bmatrix} \frac{1}{r^2}\frac{\partial^2}{\partial \phi^2} + \frac{\partial^2}{\partial z^2} & -\frac{\partial^2}{\partial z \partial r} & -\frac{1}{r^2}\frac{\partial}{\partial \phi} - \frac{1}{r}\frac{\partial^2}{\partial \phi \partial r} \\ -\frac{1}{r}\frac{\partial}{\partial z} - \frac{\partial^2}{\partial r \partial z} & \frac{\partial^2}{\partial r^2} + \frac{1}{r}\frac{\partial}{\partial r} + \frac{1}{r^2}\frac{\partial^2}{\partial \phi^2} & -\frac{1}{r}\frac{\partial^2}{\partial \phi \partial z} \\ \frac{1}{r^2}\frac{\partial}{\partial \phi} - \frac{1}{r}\frac{\partial^2}{\partial \phi \partial r} & -\frac{1}{r}\frac{\partial^2}{\partial \phi \partial z} & \frac{\partial^2}{\partial z^2} + \frac{\partial^2}{\partial r^2} + \frac{1}{r}\frac{\partial}{\partial r} - \frac{1}{r^2} \end{bmatrix} \tag{145}$$

Similar to the Cartesian case, a plane wave expansion is introduced, such that $\frac{\partial}{\partial \phi} \to ik_\phi$ and $\frac{\partial}{\partial z} \to ik_z$. To make comparison with the Cartesian expression easier, the following additional substitutions are made: $\frac{k_\phi}{r} \to k_z$, $k_z \to k_y$, $r \to x$. The extra terms, compared to the Cartesian case, are indicated in bold.

$$\overrightarrow{\overrightarrow{D}} = \begin{bmatrix} -\left(k_z^2 + k_y^2\right) & -ik_y\frac{\partial}{\partial x} & -\boldsymbol{\frac{ik_z}{x}} - ik_z\frac{\partial}{\partial x} \\ -\boldsymbol{\frac{ik_y}{x}} - ik_y\frac{\partial}{\partial x} & -k_z^2 + \frac{\partial^2}{\partial x^2} + \boldsymbol{\frac{1}{x}\frac{\partial}{\partial x}} & k_y k_z \\ \boldsymbol{\frac{ik_z}{x}} - ik_z\frac{\partial}{\partial r} & k_y k_z & -k_z^2 + \frac{\partial^2}{\partial x^2} + \boldsymbol{\frac{1}{x}\frac{\partial}{\partial x}} - \boldsymbol{\frac{1}{x^2}} \end{bmatrix} \tag{146}$$

Note that k_z is now dependent on r as $k_z(r) = k_\phi/r$. k_ϕ is the toroidal mode number.

Appendix D: Hot plasma limits when $k_{\parallel} \to 0$

In the limit where $k_{\parallel} \to 0$, the argument of the plasma dispersion function $\zeta_n \to \infty$. This means the following limits need to be determined:

$$\lim_{k_{\parallel}\to 0} \frac{1}{k_{\parallel}} Z\left(\zeta_n(k_{\parallel})\right) = \frac{v_l}{\omega + n\Omega_c} \lim_{\zeta_n\to\infty} \zeta_n Z(\zeta_n) \tag{147}$$

$$\lim_{k_{\parallel}\to 0} \frac{1}{k_{\parallel}} Z'\left(\zeta_n(k_{\parallel})\right) = \frac{v_l}{\omega + n\Omega_c} \lim_{\zeta_n\to\infty} \zeta_n Z'(\zeta_n) \tag{148}$$

$$\lim_{k_{\parallel}\to 0} \frac{1}{k_{\parallel}} \zeta_n(k_{\parallel}) Z'\left(\zeta_n(k_{\parallel})\right) = \frac{v_l}{\omega + n\Omega_c} \lim_{\zeta_n\to\infty} \zeta_n^2 Z'(\zeta_n) \tag{149}$$

The plasma dispersion function $Z(\zeta_n)$ is related to the Dawson integral $S(\zeta)$:

$$Z(\zeta) = i\sqrt{\pi} e^{-\zeta^2} - 2S(\zeta) \tag{150}$$

The asymptotic series of $S(\zeta)$ is given by:

$$S(\zeta_n) = \sum_{n=0}^{\infty} \frac{(2n-1)!!}{2^{n+1}\zeta^{2n+1}} = \frac{1}{2\zeta} + \frac{1}{4\zeta^3} + \frac{3}{8\zeta^5} + \frac{15}{16\zeta^7} + \cdots \tag{151}$$

So, the asymptotic expressions for $Z(\zeta_n)$ and $Z'(\zeta_n)$, without expanding the imaginary part, are given by:

$$Z(\zeta_n) \approx i\sqrt{\pi} e^{-\zeta_n^2} - \frac{1}{\zeta_n} - \frac{1}{2\zeta_n^3} - \frac{3}{4\zeta_n^5} - \cdots \tag{152}$$

$$Z'(\zeta_n) \approx -2i\sqrt{\pi}\zeta_n e^{-\zeta_n^2} + \frac{1}{\zeta_n^2} + \frac{3}{2\zeta_n^4} + \cdots \tag{153}$$

Using the asymptotic expressions, finding the limits is trivial. The imaginary part is always forced to zero due to the strong exponential decay. Only the real part survives in the limit.

$$\lim_{\zeta_n\to\infty} \zeta_n Z(\zeta_n) = -1 \tag{154}$$

$$\lim_{\zeta_n\to\infty} \zeta_n Z'(\zeta_n) = 0 \tag{155}$$

$$\lim_{\zeta_n\to\infty} \zeta_n^2 Z'(\zeta_n) = 1 \tag{156}$$

However, for very small $k_{\parallel}$, which means very large ζ_n, numerical problems might still occur. This can be avoided by avoiding small k_z, except for $k_z = 0$, which is well defined.

Appendix E: Hot plasma limits when $\boldsymbol{k_\perp \to 0}$

In the limit where $k_\perp \to 0$, the argument of the Bessel function $\lambda_j \to 0$. Each element of the dielectric tensor contains at least the exponent $e^{-\lambda_j}$, and one or more of the modified Bessel functions, or their derivatives. Most of the limits where $k_\perp \to 0$ show no special behavior around $k_\perp \to 0$, as the limits of the Bessel functions and the exponent behave nicely:

$$\lim_{\lambda_j \to 0} e^{-\lambda_j} = 1 \tag{157}$$

$$\lim_{\lambda_j \to 0} I_n = \begin{cases} 1 \text{ if } n = 0 \\ 0 \text{ otherwise} \end{cases} \tag{158}$$

$$\lim_{\lambda_j \to 0} I_n{}' = \begin{cases} \frac{1}{2} \text{ if } n = \pm 1 \\ 0 \text{ otherwise} \end{cases} \tag{159}$$

This already provides a way for effective evaluation of the Bessel function when $k_\perp = 0$. The only contribution that is slightly troublesome is:

$$\lim_{\lambda_j \to 0} \frac{nI_n}{\lambda_j} = \begin{cases} \frac{n}{2} \text{ if } n = \pm 1 \\ 0 \text{ otherwise} \end{cases} \tag{160}$$

At $n = 0$, the value of $\frac{nI_n}{\lambda_j}$ is forced to 0. For all non-zero n, l'Hôpital can be used to investigate the limit. In combination with the limit from (159) we find:

$$\lim_{\lambda_j \to 0} \frac{nI_n}{\lambda_j} = \lim_{\lambda_j \to 0} nI_n' = \begin{cases} \frac{n}{2} \text{ if } n = \pm 1 \\ 0 \text{ otherwise} \end{cases} \tag{161}$$

Appendix F: Heatmaps of $N - \zeta$ for other polynomial orders

$N_p = 2$

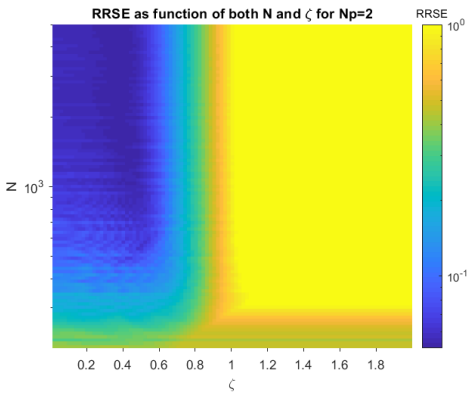

$N_p = 3$

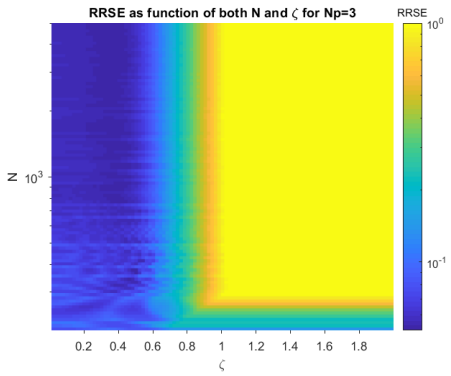

$N_p = 4$

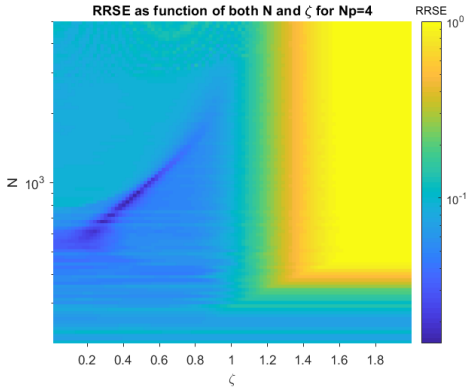

$N_p = 5$

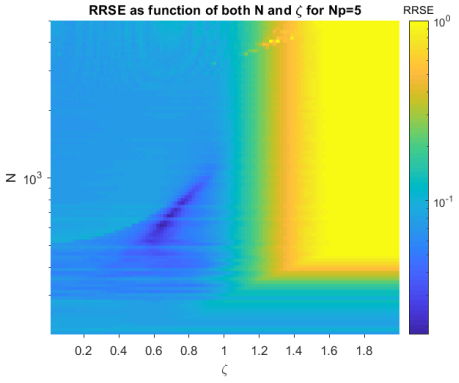

$N_p = 6$

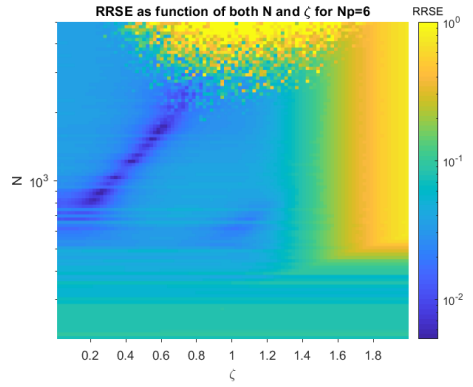

$N_p = 7$

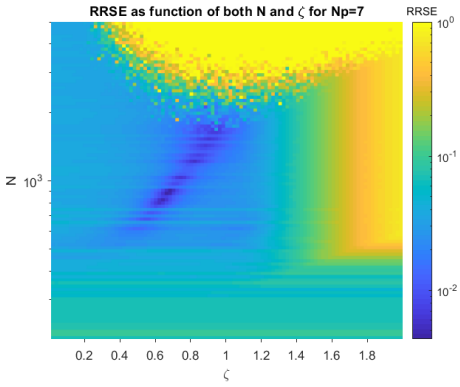

$N_p = 8$

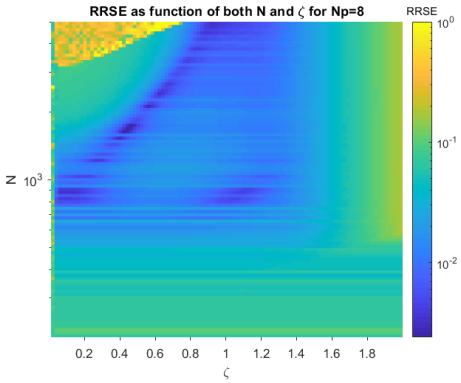

$N_p = 9$

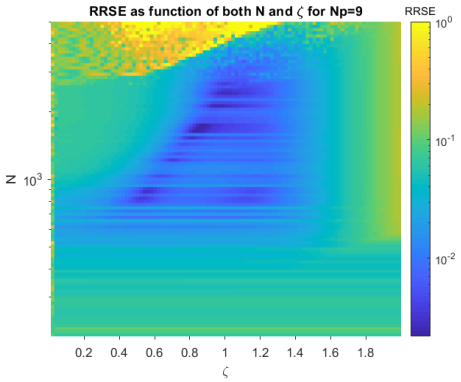

$N_p = 10$

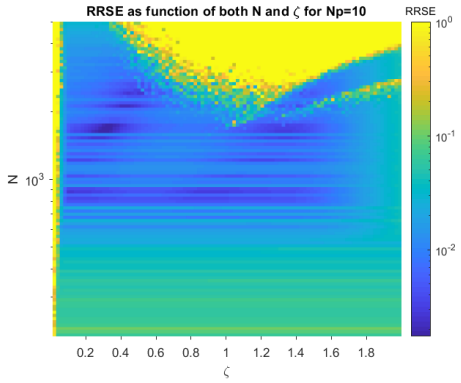

$N_p = 11$

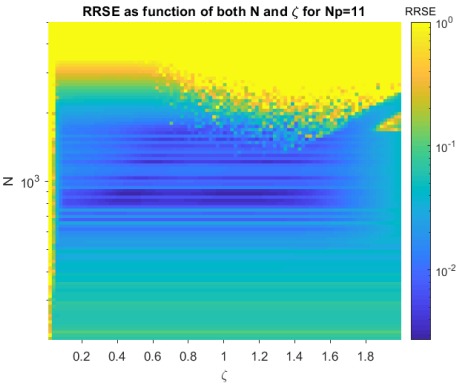

$N_p = 12$