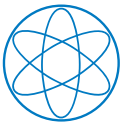

Technische Universität München

Max-Planck-Institut für Plasmaphysik

Global nonlinear fully gyrokinetic and hybrid treatments of Alfvénic instabilities in ITER

Thomas William Hayward-Schneider

Vollständiger Abdruck der von der Fakultät für Physik der Technischen Universität München zur Erlangung des akademischen Grades eines

Doktors der Naturwissenschaften

genehmigten Dissertation.

Vorsitzende:	Prof. Christine Papadakis, Ph.D.
Prüfer der Dissertation:	
	1. Prof. Dr. Sibylle Günter
	2. Prof. Laurent Villard

Die Dissertation wurde am 12.11.2019 bei der Technischen Universität München eingereicht und durch die Fakultät für Physik am 10.02.2020 angenommen.

DEPARTMENT OF PHYSICS
TECHNISCHE UNIVERSITÄT MÜNCHEN

Global nonlinear fully gyrokinetic and hybrid treatments of Alfvénic instabilities in ITER

Author: Thomas William Hayward-Schneider
Supervisor: Prof. Dr. Sibylle Günter
Advisors: Dr. Philipp Lauber
Submission Date: 12.11.2019

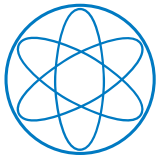

HELMHOLTZ
RESEARCH FOR GRAND CHALLENGES

EUROfusion

This thesis is dedicated to the memories of Harold Thomas Hayward (1931–1998) and Lucy Elizabeth Hayward (1985–2015).

Contents

1 Introduction

1.1 ITER as a fusion device

A large international project, the ITER tokamak [1], currently under construction in France, aims to demonstrate a significant amount of fusion power, with a goal of achieving significantly more fusion power than injected power for the first time. This will create new challenges, both for technology and physics, which we need to be able to understand and to predict in order that the project will be a success and lead to the DEMO project as the final step towards a commercial fusion reactor.

In this section, we shall introduce ITER, and tokamaks more generally, and in the next, we shall discuss the particular plasma scenario evaluated in this work, currently envisaged as the scenario in which said fusion demonstration should occur.

A tokamak is a magnetic torus built to confine a plasma using a helical magnetic field. We begin by defining some quantities relative to the torus, which can be thought of as a cylinder bent onto itself. We typically use coordinates to describe the torus, namely radial (as in the cylinder), as well as the poloidal angle (the angle in the cylinder, or the short way around the torus), and a toroidal angle (the long way around the torus). These coordinates are sketched in figure 1.1. The externally applied magnetic field is purely toroidal. In order to create the poloidal field, a current is driven through the conducting plasma, which can be viewed as the secondary winding of a transformer. This current induces a poloidal magnetic field, giving rise to the total helical magnetic field. It should also briefly be mentioned that there is another way of obtaining a helical magnetic field, namely by providing such a field externally. This approach is known as the stellarator, but is not the focus of this work. This twist of the magnetic field lines varies across the plasma, and is known as the safety factor (q). The safety factor is the number of times a magnetic field line travels around the tokamak toroidally for each poloidal transit. The profile of the safety factor across the plasma (minor) radius depends on the profile of the plasma current. To prevent certain plasma instabilities, it is normally necessary to have $q \gtrsim 1$ in the core of the plasma and $q \gtrsim 3$ at the edge of the plasma.

ITER will be the world's largest tokamak, and the first to demonstrate significant fusion power. One of ITER's design goals is to demonstrate $Q = 10$, where Q is the ratio

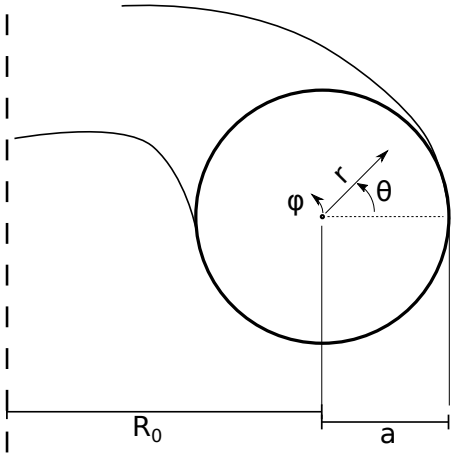

Figure 1.1: Basic coordinates in toroidal geometry. Radius r and poloidal angle θ are in the plane, whereas the toroidal angle φ is shown moving into the page. On the left, we mark the axis of symmetry. The major (R_0) and minor (a) radii are labelled.

between fusion power and externally applied heating power, eclipsing the current record of $Q \sim 0.67$ set at the Joint European Torus (JET) [2]. JET and the Tokamak Fusion Test Reactor (TFTR, a former US tokamak) [3] are, to date, the only two tokamaks which have operated with the D-T plasma mixture and the aim of demonstating the feasibility of fusion power. Of all of the various potential fusion reactions, $^2\mathrm{D} + {}^3\mathrm{T}$ has the largest fusion cross section, and is therefore the fusion reaction which ITER will pursue. The reaction in question is

$$^2\mathrm{D} + {}^3\mathrm{T} \longrightarrow {}^4\mathrm{He} + {}^1\mathrm{n}$$

where the helium ion (also known as an alpha, or α particle) takes 3.5 MeV and the neutron takes 14.1 MeV of the total 17.6 MeV released. This neutron, unaffected by magnetic fields, leaves the plasma without further interaction with the confined plasma, but the alpha particle is confined by the magnetic field, and should heat the background plasma as it slows down by collisional processes with the background. Due to the very high energy of the alpha particle with respect to the background plasma (with thermal energies up to ~ 25 keV), they have velocities comparable to, or larger than, the phase velocities of various magnetic waves. It is therefore possible for resonant interactions between these alpha particles and the waves, opening the possibility of driving instabilities, in particular Alfvén eigenmodes. Such instabilities could have a detrimental effect on the confinement of the alpha particles, reducing the efficiency of their heating the plasma, or, in the worst case, causing damage to the tokamak wall.

There can also be a positive effect, transporting or removing the slowed down helium particles (or ash) which would otherwise dilute the fusion fuel. It is therefore important to understand the interactions between alpha particles (or energetic particles (EPs) in general) and electromagnetic waves and instabilities and the consequences for the overall confined properties of a burning plasma.

1.2 The 15 MA ITER scenario

The scenario intended to demonstrate $Q = 10$ fusion in ITER is one with 15 MA of total plasma current, based upon transport simulations [4]. It features a q profile slightly below 1 at the plasma axis with a very low radial gradient (or shear, $s = \frac{r}{q}\frac{\mathrm{d}q}{\mathrm{d}r}$) in the inner region, $q_{\mathrm{edge}} \approx 4$. This profile is shown in figure 1.2b, with the detail of the inner region more clearly shown in figure 1.2c. The background plasma ions, equal densities of ^{2}D and ^{3}T are assumed to be equal in temperature, they have very flat density profiles, with the exception of the steep gradient close to the edge of the plasma, or even hollow. On the other hand, the temperature of the ions is peaked, with T_{i} at $r = 0$, $T_{\mathrm{i}}(0) \approx 21\,\mathrm{keV}$ (we typically absorb the Boltzmann constant into temperatures, such that $1\,\mathrm{eV} \approx 11.6 \times 10^3\,\mathrm{K}$, therefore this temperature is approximately $243 \times 10^6\,\mathrm{K}$). The electron temperature profile follows almost the same shape as the ion temperature profile, but is slightly larger ($T_{\mathrm{e}}(0) \approx 25\,\mathrm{keV}$), as the alpha particles primarily heat the electrons, which in turn heat the ions. The electron density profile is similarly flat, but slightly peaked. As the plasma should be macroscopically quasineutral, the difference in the densities of the ions and the electrons (where the charge on the ions $q_{\mathrm{i}} = -q_{\mathrm{e}}$) is accounted for by impurity ions in the plasma, namely beryllium (one of the materials of the plasma facing wall), helium ash, and the energetic alpha particles.

These plasma profiles are shown in figure 1.2d. We note the nominal parameters of ITER, the magnetic field strength on axis (in vacuum) $B_0 = 5.3\,\mathrm{T}$; major radius (the distance from the centre to the magnetic axis) $R_0 = 6.2\,\mathrm{m}$; minor radius (distance from the magnetic axis to the edge of the plasma on the outside of the tokamak) $a = 2\,\mathrm{m}$; aspect ratio (ratio between major and minor axis) $A = R_0/a \approx 3$ with the inverse aspect ratio $\epsilon = a/R_0 = 0.32$. The plasma is elongated, and somewhat triangular (or D-shaped), and seen in figure 1.2a. The Alfvén frequency (characteristic frequency for Alfvén waves) on axis $\omega_{\mathrm{A}} = 167.5\,\mathrm{kHz}$.

In figure 1.2e, we plot the alpha particle density profile, showing also a double alpha particle density profile. It is this artificially increased density profile which we shall primarily evaluate in this work for several reasons. Firstly, considering the double density case is chosen in line with previous work in order to allow comparison. Secondly, we wish to be pessimistic with respect to the available drive of the modes,

preferring to err on the side of increased instability, and also to consider the question of where (relative to nominal parameters) threshholds for transitions between different EP transport regimes will occur. It is only by demonstating our ability to handle such regimes that we can be able to know that this possibility was properly considered. This was the reason that some previous work also considered double density. Thirdly, we consider the possibility that the flat density profiles of this scenario are unrealistic, with some predictions suggesting also peaked density profiles as well as temperature profiles [5–7]. Given that we know that the local fusion power (the source of alpha particles) depends on the square of the deuterium and tritium density, such profiles would deliver steeper gradients in the alpha particle density. We shall later show that what drives the instabilities investigated in this work unstable is the radial gradient of the energetic alpha particle pressure.

Finally, in chapter 4, we shall also consider a modified scenario, in which we have an even flatter safety factor profile in the plasma core. This is shown as an additional line in figure 1.2b.

1.3 Observation of energetic particle driven physics on current experiments

Although current experiments do not operate with tritium, and therefore do not have fusion born alpha particles (except for the now old experiments of JET and TFTR already alluded to), they do have sources of energetic particles. These are found as a consequence of some of the plasma heating methods, in particular neutral beam injection (NBI) and ion cyclotron resonance heating (ICRH). It should be remarked that ITER will also have these heating methods as auxilliary heating to start the discharge, and for control purposes. Due to the nature of these two heating systems, the distribution functions are very different from each other, and from the alpha particle distribution function seen from fusion reactions, but can still be capable of driving instabilities.

However, as the energetic ions in most current tokamaks are sub-Alfvénic (the characteristic velocity of the waves, we shall discuss in more detail later), the drive is by a slightly different resonance mechanism. There have been a few exceptions to this, for example, JT-60U, a tokamak previously operating in Japan, had relatively high energy NBI injection energy (300 keV) compared to most machines ($\lesssim$ 100 keV). Also in spherical tokamaks, a type of tokamak with a very small aspect ratio, the beams are often super-Alfvénic, although other properties of these devices give rise to substantial differences in behaviour. As mentioned above, radial gradients in the EP pressure gradient can act as a source of free energy in the system, able to drive the instabilities via resonant interactions via inverse Landau damping, which we shall

explain in chapter 2.

Even though the modes detected by experimental observations are in nonlinear quasi-saturated states, many linear properties of the modes persist. Although direct measurements of energetic particle driven modes are difficult, signal frequency, and some spatial information can be detected. Transport due to these instabilities is rarely directly observed, but can sometimes be inferred by discrepancies between other signals. In the cases where the instabilities cause losses from the plasma, the energy and position of the ejected particles can be measured [9].

Various experimental observations of Alfvénic instabilities have demonstrated quasi-linear behaviour (where mode properties such as frequency and spatial structure are preserved from the linear modes, and experimental amplitudes can be predicted by scaling arguments) [10]; nonlinear coupling between modes (where, for example, mores are excited via wave-coupling, which has been demonstrated by performing bicoherence analysis) [11]; non-perturbative nonlinear features such as frequency chirping (where measured frequencies may repeatedly sweep up, down, or symmetrically up and down as a consequence of the formation of phase-space structures [12]); and nonlinear transport modifications, such as the enhancement of losses in the presence of multiple instabilities [13], or highly nonlinear avalanche-like events in which particles are lost as a large amount of plasma is ejected.

A particularly notable example are the so called Abrupt Large amplitude Events (ALEs) [14, 15], observed on the aforementioned JT-60U tokamak. During these short lived events, very large amplitude instabilities were observed, at amplitudes far larger than achievable without strong nonlinear coupling.

1.4 Previous modelling efforts

The problem of Alfvén eigenmodes in the ITER 15 MA scenario has been modelled extensively in the past. Even for linear physics not all models agree, with differences attributable largely to the different damping mechanisms. Models applied to the problem include initial value simulations from perturbative MHD-hybrid [16–19] and gyrokinetic-hybrid [20], non-perturbative MHD-hybrid [21], and local gyrokinetic [22] codes, and eigenvalue calculations using a linear gyrokinetic [8, 23] code. In particular for nonlinear perturbative hybrid models, all linear discrepancies are carried over to nonlinear predictions. These models will be explained in the next chapter.

On the other hand, global initial value electromagnetic fully-gyrokinetic codes have until recently been unable to address the parameter regimes needed for simulating Alfvén eigenmodes in the ITER scenario. Electromagnetic models are those in which the (parallel) fluctuations of the electromagnetic potential are solved as well as the

fluctuations in the electrostatic fields, necessary for considering finite plasma pressure and for retaining Alfvén waves. Global models are those which retain parameter variation across the simulated domain, with local models performing an expansion around a single radial position. Here, we define fully-gyrokinetic as treating all the plasma species, including electrons, kinetically. Previous works in simpler (with respect to plasma parameters such as ϵ, β or ρ^*, see definition below) scenarios include references [24–29]. These models are also capable of (and original designed for) studying electrostatic and electromagnetic turbulence and anomalous transport.

The difficulty in addressing the ITER 15 MA scenario with global gyrokinetic models is related to the large size, tight aspect ratio, and high pressure. The expected plasma parameters of ITER far exceed the parameter space occupied by present machines. The relevant quantities, present in the equations solved, normalize the physical quantities. For example, the ratio between the plasma pressure to the magnetic pressure ($\beta = \frac{nT}{B^2/2\mu_0}$, which can be evaluated for a particular plasma species or for the total plasma); and the ratio between the equilibrium scale and the ion sound Larmor radius ($\rho^* = \rho_s/a$, we shall explain the Larmor radius in the next chapter). Taking parameters for the plasma core, we obtain $\beta_{\text{thermal}} \sim 6.8\%$, $\beta_\alpha \sim 1.15\%$, and $\rho^* \sim 1/470$.

Finally, we briefly mention relevant modelling done neither of the ITER 15 MA scenario, nor with fully-gyrokinetic models, but exhibiting strong nonlinear or non-perturbative features. These include demonstrations of frequency chirping of Alfvén modes [30], nonlinear enhancement of losses [31], and simulations of the previously mentioned nonlinear bursting ALEs [32, 33].

In this work, we go beyond the previous work by applying a more complete model (which forms a superset of the various different models previously used) to the ITER 15 MA scenario.

1.5 Aims and Structure of this thesis

The scope of this work is as follows. Given the importance of predictive capabilities for alpha particle physics in ITER, one of the issues which we cannot directly study on present day experiments, and the discrepancies found linearly and nonlinearly by various reduced models, in this work we exploit a hierarchy of models to gain the understanding necessary to tackle this important issue with the most advanced model available. These results can then inform the community about the expected behaviour, and can also be used to help to understand when reduced models will be applicable.

The structure of this thesis is as follows. In chapter 2, a theoretical background of the important physical mechanisms and how they can be treated is introduced, before a

more detailed description of the models implemented in the numerical tools, and their features is presented in chapter 3.

Then we present the findings of this work, with linear and nonlinear results from perturbative modelling presented in chapter 4, the first application of a global gyrokinetic model to this question in chapter 5, with nonlinear results in chapter 6, we are able to make predictions based on nonlinear global non-pertubative gyrokinetic results.

Finally, we conclude the thesis with a summary of the results and major achievements of this work in chapter 7, and a discussion about the missing pieces to this story, open issues, and how they can be addressed.

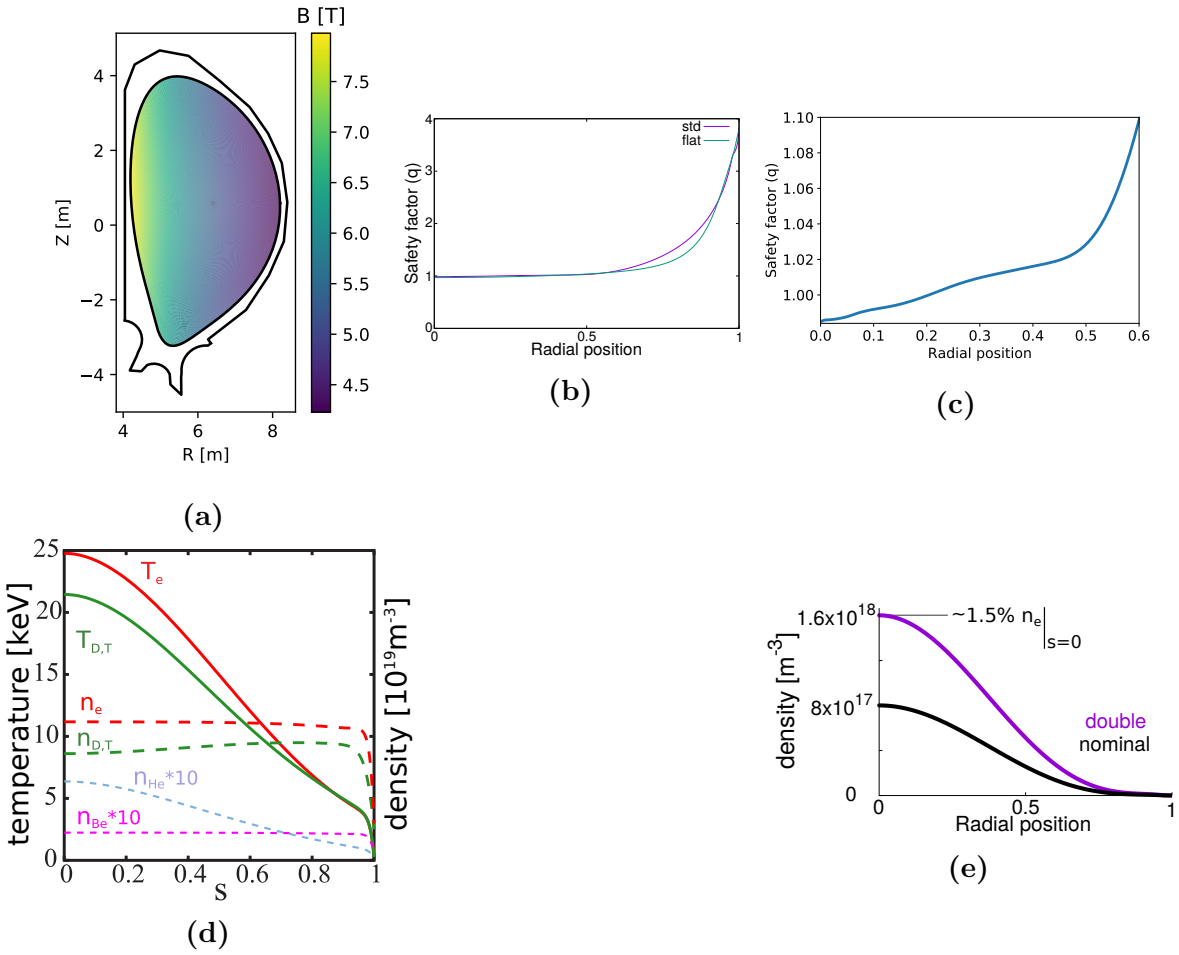

Figure 1.2: Scenario parameters for the ITER 15 MA scenario. We show: in (a) the outline of the equilibrium and the magnetic field strength; (b) and (c) the safety factor profile across the full radius and a zoom of the inner part of the radius respectively, in (b) we also plot the safety factor of an alternate 'flat' profile; (d) plasma density and temperature profiles for the background thermal plasma species (temperatures are shown with solid lines and densities with dashed lines); and (e) the alpha particle density profile, showing both nominal and double parameters. The image in (d) is modified from reference [8]. The radial coordinate in figures (b-e) is the square root of the normalized poloidal flux ($s = \sqrt{\psi/\psi_{\text{edge}}}$), and ψ is the poloidal magnetic flux.

2 Theory

In this chapter, we introduce the theoretical properties and methods used in this work. We do not attempt to provide a complete introduction, or formal derivations for the methods used, for these are better served by existing references, for example references [34–38]. We aim, however, to provide an overview.

2.1 Particle motion in a magnetic field

In the core of fusion plasmas, we deal with fully-ionized gases following strong magnetic guide-fields. The most important behaviour for us to consider is therefore the Lorentz force, perpendicular to the velocity of a particle and to the magnetic field (we neglect for now the contribution from an electric field),

$$\boldsymbol{F} = q\boldsymbol{v} \times \boldsymbol{B} \tag{2.1}$$

Separating the motion parallel to the magnetic field from that perpendicular to the magnetic field, we see that the perpendicular motion of a particle circles (or gyrates) around a magnetic field line. The frequency associated with such motion is the particle gyrofrequency,

$$\omega_\mathrm{c} = |q|\, B/m \tag{2.2}$$

and for ions in the ITER 15 MA scenario is ≈ 30 MHz, and the Larmor radius associated with this gyration depends on the perpendicular component of the velocity,

$$\rho_L = v_\perp/\omega_c \tag{2.3}$$

As the force acts perpendicular to the motion of the particle it performs no work, and we separate the fast gyration motion from the slower motion of the centre of the gyration. The motion of the centre of this helical motion is known as the guiding centre motion, and it approximately follows the magnetic field. The deviations of the motion away being parallel to the magnetic field are know as drifts.

In the case of a general drift, due to a force $\boldsymbol{F}$ applied to a particle, the drift velocity

is therefore,

$$\boldsymbol{v}_{\text{drift}} = \boldsymbol{F} \times \boldsymbol{B}/(qB^2) \tag{2.4}$$

noting as a result of the dependence on the charge q, that forces which are independent on the sign of the charge lead to opposite drifts for electrons and ions.

We know that the adiabatic invariant,

$$\mu = mv_{\perp}^2/(2B) \tag{2.5}$$

is conserved. This tells us that as the total field strength changes, the perpendicular velocity must change to conserve μ. However, as the total kinetic energy of the particle in a static magnetic field should also be conserved, any change to the perpendicular velocity must also change the parallel velocity. For sufficient changes in the magnetic field strength, relative to the initial parallel and perpendicular components of a particle's velocity, the particle may have its parallel motion reversed. This leads us to the concept of the magnetic mirror, or particle trapping.

Since in a tokamak, the particles follow the helical magnetic field lines, they transit from the inside to the outside of the tokamak (i.e. changing R at approximately constant r). We also know that the magnetic field strength in a tokamak decays with the radius from the symmetry axis ($B \sim 1/R$), a particle can become trapped (or mirrored) when following a field line.

2.2 Wave-particle resonant interaction

Waves propagate at a phase velocity $v_{\text{phase}} = \omega/k$. Any particle also moving around this velocity will see a constant phase of the wave. Therefore, any forces of the wave felt by the particle will be felt for enough time for the particle to become accelerated or deccelerated. We discuss this process in the case of Landau damping [39] in a one-dimensional system [40].

The force experienced due the potential of a wave is the electric field $E(x) = -\nabla\phi(x)$. In the frame of the co-moving frame of the wave, this field is constant. Considering three particles initially centered in the well of a small electric field, with velocities $v_{\text{phase}} - \epsilon$, v_{phase}, and $v_{\text{phase}} + \epsilon$ (where ϵ indicates a small variation). The particle with $v = v_{\text{phase}}$ is not accelerated by the wave. The particle with $v = v_{\text{phase}} - \epsilon$ initially moves in the negative x direction in wave frame. This particle then experiences a force due to the electric field in the positive direction, gaining energy from the wave. The particle with $v = v_{\text{phase}} + \epsilon$ initially moves in front of the wave, and is deccelerated by the electric field, transferring energy from the particle to the wave. Therefore, the net energy transfer depends on the gradient of the distribution function locally around the

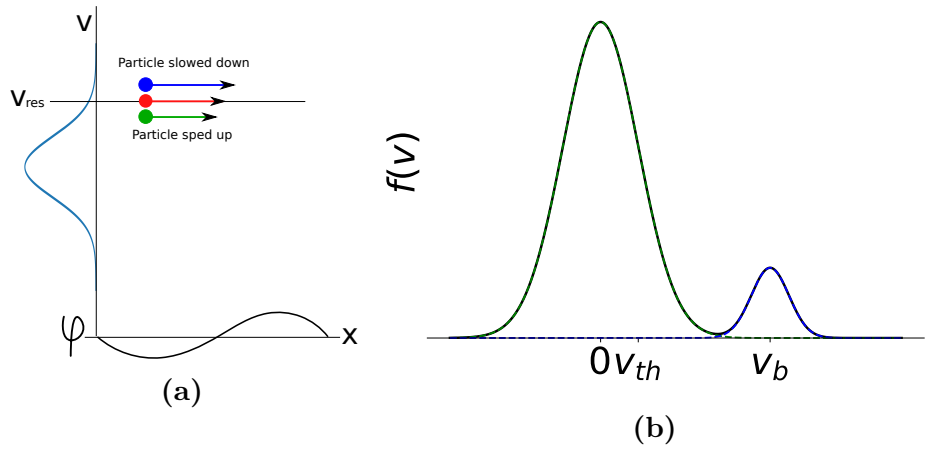

Figure 2.1

Figure 2.2: Left: a simple representation of Landau damping. Three particles (green, red, and blue in order of increasing velocity) are initially in the centre of the wave's potential well (see bottom axis). Net transfer depends on the gradient of the distribution function (see left axis). Right: Bump-on-tail distribution. Unlike the Maxwellian distribution (left), this distribution is unstable over a range of wave-numbers.

phase velocity. In the typical case of a Maxwellian distribution, the gradient of the distribution is negative for all (positive) values of v. Therefore such a mechanism will always damp the pertubation. This effect is known as Landau damping.

However, in the case where the distribution function has a reversal, such as that in the bump-on-tail distribution shown in figure 2.1b, there exists a range over which a wave's phase velocity would lead resonant drive of the pertubation (through the same mechanism as was present in Landau damping). This effect is known as inverse Landau damping, and corresponds to a linearly unstable system.

2.3 Modelling collective plasma behaviour

To model a plasma by considering the dynamics of every particle, for which we might have $N \sim 10^{22}$–10^{23} particles in a fusion reactor, is impossible. Therefore, we instead consider a distribution function as our fundamental representation, from which other quantities can be extracted by evaluating moments.

At our most general level, for a fully-ionized, confined plasma, we assume that the behaviour of a plasma can be described by the collisionless non-relativistic Vlasov equation [41], describing the evolution of the six-dimensional phase space of a plasma

species s,

$$\frac{\partial f_s}{\partial t} + \boldsymbol{v} \cdot \nabla f_s + \frac{q_s}{m_s}(\boldsymbol{E} + \boldsymbol{v} \times \boldsymbol{B}) \cdot \frac{\partial f_s}{\partial \boldsymbol{v}} = 0 \tag{2.6}$$

where the right hand side can have a collision operator or sources or sinks added as appropriate.

This must be coupled to a the Maxwell equations for solving the electromagnetic fields.

$$\nabla \times \boldsymbol{B} = \mu_0 \left(\boldsymbol{J} + \epsilon_0 \frac{\partial \boldsymbol{E}}{\partial t} \right) \tag{2.7}$$

$$\nabla \times \boldsymbol{E} = -\frac{\partial \boldsymbol{B}}{\partial t} \tag{2.8}$$

$$\nabla \cdot \boldsymbol{E} = \frac{\rho}{\epsilon_0} \tag{2.9}$$

$$\nabla \cdot \boldsymbol{B} = 0 \tag{2.10}$$

where the charge density (ρ) and current density (J) are themselves built from the distribution function f.

These equations, six-dimensional, and retaining the fast gyration of the particles around a magnetic field, are impractical to solve for fusion plasmas (although are indeed used for astrophysical plasmas). For our problem, we need to use a theoretical framework in which the fast time scales have been removed, and in which the dimensionality has been reduced.

2.3.1 Gyrokinetics

The first such method is to take the Vlasov-Maxwell system of equations in the presence of a strong magnetic field, and to recast it in terms of variables in which, by ordering, we can ignore the gyro-phase. This allows us to remove the fast gyration from the system as well as reducing the dimensionality of phase space from six-dimensional to five-dimensional. This process yields the gyrokinetic model [42], traditionally performed by Lie transforms (but which can also be performed by an averaging method [43, 44]) amounts to an averaging of the particle motion around the gyromotion, giving the equations in terms of a gyrocenter, parallel velocity (or momentum) and magnetic moment (representing the perpendicular velocity as described above). We can therefore think of a series of slowly moving charge rings, instead of point charges. Below, we discuss the models of gyrokinetics, as well as discuss the difficulties of electromagnetic simulations (the cancellation problem as we shall see below), following the work of [45].

When deriving equations for use in numerical models in tokamaks, we must perform

an expansion in terms of small parameters [45, 46], the two most important of which correspond to the ratio between the thermal Larmor radius and the equilibrium variation scale ($\epsilon_\mathrm{B} = \rho_\mathrm{th}/L_\mathrm{B}$, where $L_B = |\nabla B/B|^{-1}$); and to the relative electromagnetic fluctuation amplitudes ($\epsilon_\delta = \omega/\omega_\mathrm{c} \sim k_\parallel/k_\perp \sim q\delta\phi/T \sim \delta B/B \ll 1$)

We assume that the perturbed electromagnetic potential is purely parallel to the magnetic field, setting $A_\perp = 0$, and by constructing a perturbed Poincaré-Cartan form [42, 45],

$$\gamma = q\boldsymbol{A}^*(\boldsymbol{R}) \cdot \mathrm{d}\boldsymbol{R} + \frac{m}{q}\mu \mathrm{d}\theta - \left(\frac{mv_\parallel^2}{2} + \mu B\right)\mathrm{d}t + q\left[A_\parallel(\boldsymbol{x})\boldsymbol{b}\cdot\mathrm{d}\boldsymbol{x} - \phi(\boldsymbol{x})\mathrm{d}t\right] \tag{2.11}$$

we observe that as $\boldsymbol{x} = \boldsymbol{R} + \boldsymbol{\rho}(\theta)$, the perturbed fields depend on the gyro-phase θ, with $\boldsymbol{R}$ the guiding-centre position, and ρ the gyroradius, $\boldsymbol{B}$ the background magnetic field, $\boldsymbol{A}^* = \boldsymbol{A} + (mv_\parallel/q)\boldsymbol{b}$ is the extended magnetic potential, $\boldsymbol{A}$ is the magnetic potential corresponding to $\boldsymbol{B}$, and $\boldsymbol{b} = \boldsymbol{B}/B$ is the unit vector of the background magnetic field.

By using the aforementioned Lie transform methods $\Gamma = \mathrm{e}^{\hat{G}}\gamma + \mathrm{d}S$ (as detailed in reference [42]), there is some freedom as to how to treat the perturbed magnetic potential. The two standard methods are the $p_\parallel$ (Hamiltonian) formulation,

$$\Gamma = q\boldsymbol{A}^*\cdot\mathrm{d}\boldsymbol{R} + \frac{B}{\Omega}\mu\mathrm{d}\theta - \left(\frac{mp_\parallel^2}{2} + \mu B + q\langle\phi - p_\parallel A_\parallel\rangle\right)\mathrm{d}t \tag{2.12}$$

and the $v_\parallel$ (symplectic) formulation,

$$\Gamma = q\boldsymbol{A}^*\cdot\mathrm{d}\boldsymbol{R} + \frac{B}{\Omega}\mu\mathrm{d}\theta + q\langle A_\parallel\rangle\boldsymbol{b}\cdot\mathrm{d}\boldsymbol{R} - \left(\frac{mv_\parallel^2}{2} + \mu B + q\langle\phi\rangle\right)dt \tag{2.13}$$

In the $v_\parallel$ formulation, a partial derivative $\partial\langle A_\parallel\rangle/\partial t$ appears in the equation for $\dot{v}_\parallel$, requiring implicit solvers (impractical for Lagrangian methods). On the other hand, the $p_\parallel$ formulation has no such problem in the equations for the trajectories. However, in the $p_\parallel$ formulation, the field equations take the following form,

$$\sum_{s\neq e}\int \frac{q_s^2 F_{0s}}{T_s}\left(\phi - \langle\phi\rangle\right)\delta_\mathrm{gy}\mathrm{d}^6Z = \sum_s q_s n_{1s} \tag{2.14}$$

$$\sum_s \frac{\beta_s}{\rho_s^2}\langle\overline{A_\parallel}\rangle_s - \nabla_\perp^2 A_\parallel = \mu_0\sum_s j_{\parallel 1s} \tag{2.15}$$

noting that additional skin terms appear in the Ampère's law (equation 2.15) propor-

tional to β_e/ρ_e^2 and β_i/ρ_i^2. As the electron mass features, $A_\parallel \beta_e/\rho_e^2 = \mu_0 n_0 e^2 A_\parallel/m_e$, these terms (which are unphysical) can be very large, and must be accurately cancelled by the (also unphysical) contribution of the parallel current.

This is the origin of the so-called "cancellation problem" in gyrokinetics [47–49], especially limiting global fully kinetic electromagnetic particle-in-cell codes. Although a "control variate" method, based upon an iterative scheme [46, 50] made simulations possible, recent methods [45, 51–54], based on a new formulation, the mixed-variables formulation offer a significant improvement over the "control variate" method. It is this mixed-variables formulation that is used for the gyrokinetic simulations with ORB5 in this work. The equations solved are detailed in chapter 3.

2.3.2 Magnetohydrodynamics (MHD)

In the above section, we have shown how one can arrive at a kinetic model using an ordering which makes it suitable for fusion parameters. Another method is to reduce the dimensionality of the problem by taking moments of the distribution function, and casting the equations in terms of these moments. This will allow us to arrive [55] at a fluid model for the behaviour of the plasma.

By writing macroscopic quatities in terms of moments of the distribution function, e.g. particle number density,

$$n_s(\boldsymbol{r}, t) = \int f_s(\boldsymbol{r}, \boldsymbol{v}, t) \mathrm{d}^3\boldsymbol{v}$$

flow velocity,

$$\boldsymbol{u}_s(\boldsymbol{r}, t) = \frac{1}{n_s(\boldsymbol{r}, t)} \int \boldsymbol{v} f_s(\boldsymbol{r}, \boldsymbol{v}, t) \mathrm{d}^3\boldsymbol{v}$$

and the sources for Maxwell's equations, charge density,

$$\rho_c = \sum_s e_s n_s$$

and current density,

$$\boldsymbol{j} = \sum_s e_s n_s \boldsymbol{u_s}$$

We note, however, that moments of order k are coupled to those of order $k-1$ and $k+1$, thus creating the problem of closing the model. To do this, we require a closure, using the adiabatic equations of state.

2.4 Alfvén waves and eigenmodes

Starting from the MHD equations, and following reference [56],

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \boldsymbol{v}) = 0$$

$$\rho \frac{\partial \boldsymbol{v}}{\partial t} + \rho \boldsymbol{v} \cdot \nabla \boldsymbol{v} + \nabla P - \boldsymbol{j} \times \boldsymbol{B} = 0$$

$$\frac{\partial P}{\partial t} + \boldsymbol{v} \cdot \nabla P + \gamma P \nabla \cdot \boldsymbol{v} = 0$$

$$\boldsymbol{E} + \boldsymbol{v} \times \boldsymbol{B} = 0$$

$$\nabla \cdot \boldsymbol{B} = 0$$

$$\nabla \times \boldsymbol{B} - \mu_0 \boldsymbol{j} = 0$$

$$\frac{\partial \boldsymbol{B}}{\partial t} + \nabla \times \boldsymbol{E} = 0$$

(where P is the pressure, and γ is the adiabatic index), we split quantities into a constant background part and a small time varying part, $Q(\boldsymbol{r}, t) = Q_0(\boldsymbol{r}) + Q_1(\boldsymbol{r}, t)$, assuming that $Q_1 \ll Q_0$. We therefore linearize the MHD equations, assuming also that $v_0 = 0$.

$$\frac{\partial p_1}{\partial t} + \nabla \cdot (\rho_0 \boldsymbol{v}_1) = 0$$

$$\rho_0 \frac{\partial \boldsymbol{v}_1}{\partial t} + \nabla P_1 - \boldsymbol{j}_1 \times \boldsymbol{B}_0 - \boldsymbol{j}_0 \times \boldsymbol{B}_1 = 0$$

$$\frac{\partial P_1}{\partial t} + \boldsymbol{v}_1 \cdot \nabla P_0 + \gamma P_0 \nabla \cdot \boldsymbol{v}_1 = 0$$

$$\nabla \cdot \boldsymbol{B}_1 = 0$$

$$\frac{\partial \boldsymbol{B}_1}{\partial t} + \nabla \times \boldsymbol{E}_1 = 0$$

$$\nabla \times \boldsymbol{B}_1 - \mu_0 \boldsymbol{j}_1 = 0$$

$$\boldsymbol{E}_1 + \boldsymbol{v}_1 \times \boldsymbol{B}_0 = 0$$

We introduce a displacement vector $\boldsymbol{\xi}(\boldsymbol{r}, t)$, such that $\boldsymbol{v}_1 = \frac{\partial \boldsymbol{\xi}}{\partial t}$. By elimination of $\boldsymbol{j}_1$ and $\boldsymbol{E}_1$ from the above equations, we arrive at,

$$\rho_0 \frac{\partial^2 \boldsymbol{\xi}}{\partial t^2} + \nabla P_1 + \frac{1}{\mu_0} \Big(\boldsymbol{B}_1 \times (\nabla \times \boldsymbol{B}_0) + \boldsymbol{B}_0 \times (\nabla \times \boldsymbol{B}_1) \Big) = 0$$

$$P_1 + \boldsymbol{\xi} \cdot \nabla P_0 + \gamma P_0 \nabla \cdot \boldsymbol{\xi} = 0$$
$$\boldsymbol{B}_1 + \nabla \times (\boldsymbol{B}_0 \times \boldsymbol{\xi}) = 0$$

By eliminating P_1 and $\boldsymbol{B}_1$, we arrive at the force-operator equation of linearized MHD,

$$\rho_0 \frac{\partial^2 \boldsymbol{\xi}}{\partial t^2} = \boldsymbol{F}(\boldsymbol{\xi}) \tag{2.16}$$

with

$$\boldsymbol{F}(\boldsymbol{\xi}) \equiv \nabla(\boldsymbol{\xi} \cdot \nabla P_0 + \gamma P_0 \nabla \cdot \boldsymbol{\xi}) + \frac{1}{\mu}\left[\Big(\nabla \times \nabla \times (\boldsymbol{\xi} \times \boldsymbol{B}_0)\Big) \times \boldsymbol{B}_0 + (\nabla \times \boldsymbol{B}_0) \times \Big(\nabla \times (\boldsymbol{\xi} \times \boldsymbol{B}_0)\Big)\right]$$

Assuming that the time dependent quantities can be written as $Q_1(\boldsymbol{r}, t) = Q_1(\boldsymbol{r})\mathrm{e}^{-i\omega t}$, we can write the force-operator equation as an eigenvalue problem,

$$-\omega^2 \boldsymbol{\xi} = \frac{1}{\rho_0} \boldsymbol{F}(\boldsymbol{\xi}) \tag{2.17}$$

As $\boldsymbol{F}(\boldsymbol{\xi})$ is Hermitian [57], the eigenvalues ω^2 are real, and therefore the solutions ω are purely real or imaginary, growing exponentially or oscillating at constant amplitude. It also follows that the discrete modes forming the eigenvectors are orthogonal. In the case of a purely real frequency ω, the addition of other effects (e.g. an energetic particle species) may modify the MHD eigenvalues, giving weakly growing or weakly damped modes.

In the limiting case of a homogeneous plasma, we can perform the analysis of the solutions to the MHD wave equation. By Fourier transform of the operator, we arrive at,

$$\omega^2 \rho_0 \boldsymbol{\xi} + \gamma P_0 (\boldsymbol{k} \cdot \boldsymbol{\xi}) \boldsymbol{k} + \frac{1}{\mu_0}\left(\boldsymbol{k} \times \Big[\boldsymbol{k} \times (\boldsymbol{\xi} \times \boldsymbol{B}_0)\Big]\right) \times \boldsymbol{B}_0 = 0 \tag{2.18}$$

with $\boldsymbol{k}$ the wave vector.

Aligning $\boldsymbol{B}_0 = B_0 \hat{\boldsymbol{z}}$ and $\boldsymbol{k} = k_\perp \hat{\boldsymbol{y}} + k_\parallel \hat{\boldsymbol{z}}$, we arrive at the matrix equation,

$$\begin{pmatrix} \omega^2 - k_\perp^2 v_A^2 & 0 & 0 \\ 0 & \omega^2 - k_\perp^2 v_s^2 - k^2 v_A^2 & -k_\perp k_\parallel v_s^2 \\ 0 & -k_\perp k_\parallel v_s^2 & \omega^2 - k_\parallel^2 v_s^2 \end{pmatrix} \begin{pmatrix} \xi_x \\ \xi_y \\ \xi_z \end{pmatrix} = 0 \tag{2.19}$$

with $v_A = \sqrt{B_0^2 / \mu_0 \rho_0}$ the Alfvén speed and $v_s = \sqrt{\gamma P_0 / \rho_0}$ the adiabatic sound speed.

Finding solutions by setting the determinate of the matrix to zero,

$$(\omega^2 - k_\parallel^2 v_A^2)\left[\omega^4 - (v_s^2 + v_A^2) k^2 \omega^2 + (k k_\parallel v_s v_A)^2\right] = 0 \tag{2.20}$$

we find this affords three solutions. The first solution forms the Alfvén branch, $\omega^2 = k_\parallel^2 v_A^2$, independent of $k_\perp$, incompressible, with $\boldsymbol{v}_1$ and $\boldsymbol{B}_1$ both perpendicular to $\boldsymbol{B}_0$. The other two branches include coupling between the sound and Alfvén compressions, the fast and slow magnetosonic waves. Due to the stronger coupling between energetic particles and the shear Alfvén wave in tokamak plasmas, we consider only this first branch, the shear Alfvén wave (SAW).

Although in a homogeneous plasma, the waves oscillate without dispersion, in an inhomogeneous plasma, however, the dispersion relation depends on position, $\omega(r) = k_\parallel(r)v_A(r)$, and therefore a wave packet will disperse, a property known as phase mixing. This effective damping is proportional to the radial derivative of the dispersion relation,

$$\gamma_d \sim -\left|\frac{\mathrm{d}}{\mathrm{d}r}k_\parallel(r)v_A(r)\right| \tag{2.21}$$

In a periodic cylindrical system, we can apply a decomposition in radius, and poloidal, and z-direction wavenumbers.

$$\boldsymbol{\xi}(r,\theta,z,t) = \sum_m \xi_m(r)\mathrm{e}^{i(m\theta+k_z z+\omega t)}$$

where m is the poloidal mode number. In the case that this cylinder is bent into a torus, we add now also periodicity in the toroidal (previously z) direction, introducing the toroidal mode number n.

$$\boldsymbol{\xi}(r,\theta,\varphi,t) = \sum_n \sum_m \xi_{n,m}(r)\mathrm{e}^{i(m\theta+n\varphi+\omega t)}$$

However, due to the toroidal bending, and because the equilibrium magnetic field strength is now a function of the poloidal angle, we find that the poloidal harmonics m, previously independent, are now coupled to $m \pm 1$, with the strength of the coupling proportional to the inverse aspect ratio $\epsilon = a/R_0$.

Now, at the position where the continuum spectrum for m and $m+1$ should have crossed according to the expressions for the cylinder, due to the coupling, we see a gap appear in the spectrum. This gap is known as the toroidicity-induced Alfvén eigenmode [58] (TAE, also referred to as the toroidal Alfvén eigenmode). We find this position by identifying the points at which the wave vectors of m and $m+1$ are opposite,

$$k_{\parallel m} = \frac{1}{R}\left(n - \frac{m}{q(r)}\right) \tag{2.22}$$

$$k_{\parallel m+1} = \frac{1}{R}\left(n - \frac{m+1}{q(r)}\right) \tag{2.23}$$

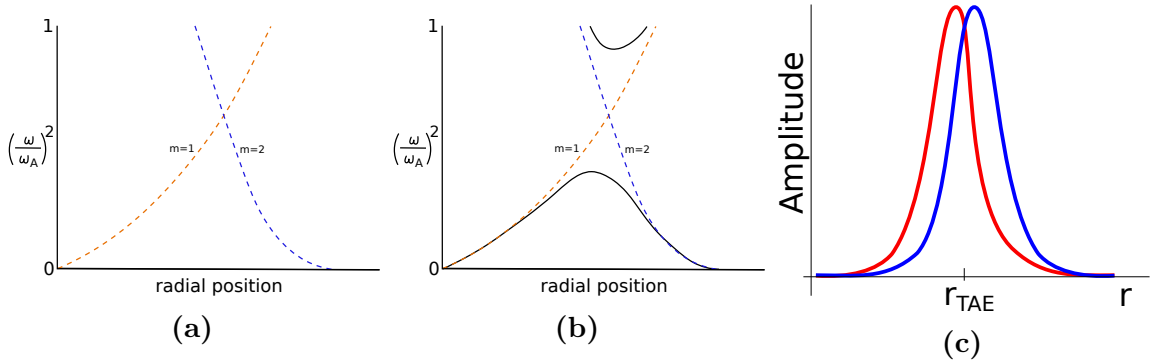

Figure 2.3: (a) The continuum spectrum for m and $m+1$ for a cylinder. (b) In black the continuum spectrum due to coupling of poloidal harmonics due to toroidicity, with the TAE gap visible. (c) A sketch of poloidal harmonics m and $m+1$ of a TAE (red and blue respectively). (a) and (b) are based on figure 1 in reference [58].

by setting $k_{\|m} - k_{\|m+1} = 0$, we arrive at,

$$q(r) = \frac{2m+1}{2n} = \frac{m+1/2}{n} \tag{2.24}$$

noting that the TAE is not located at a rational surface for n, but rather between rational surfaces.

Substituting the expressions for q and $k_{\|}$ back to the dispersion relation, we recover the TAE frequency,

$$\omega_{\mathrm{TAE}} = \frac{\omega_{\mathrm{A}}}{2q}$$

where $\omega_{\mathrm{A}} = v_A/R$.

Although the TAE is the most well known of such gap Alfvén eigenmodes, by analogy, asymmetries in the equilibrium of higher order (e.g. ellipticity, triangularity) create coupling between other pairs of harmonics, (e.g. $(m, m+2)$, $(m, m+3)$). These gaps occur at higher frequency, and are known as the ellipticity-induced Alfvén eigenmode (EAE) [59–61], triangularity-induced Alfvén eigenmodes (NAE, as its also known as the non-circularity induced Alfvén eigenmode). Gaps of higher coupling exist, but are typically of little importance and are not named.

As these modes are oscillatory eigenmodes of ideal MHD, we cannot calculate the mode drive and damping without moving to a kinetic framework. This can be gyrokinetics, which has been shown to include MHD physics [62], or a hybrid treatment, in which the modes are defined by MHD, and only the energetic particles are added for their kinetic effects. Hybrid treatments include perturbative [63–68], in which the MHD eigenvalue problem is solved, and this eigenvalue and eigenfunction is added to a hybrid code, or

non-perturbative, in which the MHD equations are solved throughout the simulation, with a contribution from the kinetically modelled EPs [69–71].

2.4.1 TAE drive and damping

As discussed, we need to move to a kinetic treatment in order to include the driving and damping mechanisms. For the mode drive, as we discussed for the case of (inverse) Landau damping, only the particles which see approximately constant phase are able to exchange energy with the mode. Therefore the resonance condition for a particle is that the the phase Θ should on average be constant.

$$\langle \dot{\Theta} \rangle \approx 0$$

We know that passing particles orbit the plasma approximately according to the field lines. However, they deviate from this with a width,

$$\Delta_{b,\alpha} = \left(v_{\parallel}^2 + \frac{v_{\perp}^2}{2} \right) \frac{q}{v_{\parallel} \omega_{B,\alpha}}$$

expanding this, we find the resonance becomes,

$$\frac{n}{R} v_{\parallel} - \frac{m+l}{qR} v_{\parallel} - \omega = 0$$

which gives us the parallel resonant velocity,

$$v_{\parallel} = \frac{qR\omega}{nq - m - l}$$

Assuming that a TAE the mode is primarily localized in the region around $q = \frac{m+1/2}{n}$, then we find

$$v_{\parallel} = \frac{v_A}{|2l - 1|}$$

The main resonance occuring for $l = 0$, namely $v_{\parallel} = v_A$ with sideband resonances at $v_{\parallel} = v_A/3$ the only resonances available on machines with sub-Alfvénic particles.

For an isotropic alpha particle distribution, reference [72] shows that the growth rate comes from the radial gradient of the distribution function,

$$\gamma/\omega_0 \approx \frac{9}{4} \left[\beta_\alpha \left(\frac{\omega_{*,\alpha}}{\omega_0} - \frac{1}{2} \right) F - \beta_e \frac{v_A}{v_e} \right]$$

where $\omega_{*,\alpha} = -k_\theta \frac{T_\alpha}{e_\alpha B_0} \frac{\mathrm{d} \ln \beta_f}{\mathrm{d}r}$

In addition to the mode drive, there are also additional mode damping mechanisms [35]. These include the coupling of the TAE to small-scale kinetic Alfvén waves which propagate radially, thus carrying energy from the mode in a process known as radiative damping [73]; Landau damping of the electrons (primarily the trapped electrons as discussed in reference [74]) which occurs away from the TAE position on rational surfaces; Landau damping of the ions [61] (primarily the passing ions in the tail of the distribution); continuum damping [75], the aforementioned phase mixing from the radial derivative of the dispersion relation; and collisional trapped electron damping [76, 77], an effect which is neglected in this work as the collision frequency is very small.

3 Numerical tools

In order to study the physics of Alfvén eigenmodes in a future experiment, we need to use theoretical tools and numerical models. The assumptions of such models will vary, and we can create a hierarchy of different models for studying different phenomena at different levels of fidelity, with different speed and cost. Firstly, we know that we need to use electromagnetic models, necessary for the study of Alfvénic physics. We shall also want to represent either the entire plasma or at least the energetic particles using a kinetic model, in order to capture the resonant interactions between Alfvén modes and EPs.

3.1 Nonlinear perturbative drift-kinetic hybrid code HAGIS

HAGIS is a nonlinear perturbative model for evolving a system of energetic particles and linear eigenfunctions using a wave-particle interaction model. It is formulated by creating a separation of the fast (mode frequency, particle motion) and slow (amplitude, phase evolution) time scales. Only the wave-particle nonlinearity is retained. Here, we outline the model used in the HAGIS code [56], closely following reference [65].

The starting point is the total system Lagrangian

$$\mathcal{L}_{sys} = \mathcal{L}_{EP} + \mathcal{L}_{int} + \mathcal{L}_{bulk} + \mathcal{L}_{em} \tag{3.1}$$

where the terms with subscripts EP, int, bulk, and em correspond to the motion of the energetic particles in the equilibrium, the effect of the Alfvén waves on the particle motion, the background plasma contribution to the Alfvén waves, and the electromagnetic component of the Alfvén waves. The first two terms sum to give the particle Lagrangian, the last two to give the wave Lagrangian $\mathcal{L}_w$ and the last three terms together describe the wave equations.

The model is cast in terms of spatial coordinates (ψ, θ, ζ), corresponding to the poloidal flux (the radial coordinate), generalized poloidal angle, and toroidal angle respectively.

For a single particle, the guiding centre Lagrangian can be written as

$$L = e\boldsymbol{A}^* \cdot \dot{\boldsymbol{x}} + \left(\frac{m}{e}\right)\mu\dot{\xi} - \mathcal{H} \tag{3.2}$$

with the particle Hamiltonian,

$$\mathcal{H} = \frac{1}{2}mv_\parallel^2 + \mu B + e\phi \tag{3.3}$$

with $\phi(\boldsymbol{x}, t)$ the electrostatic potential and ξ the gyrophase.

Writing the perturbed electromagnetic potential,

$$\widetilde{\boldsymbol{A}}(\boldsymbol{x}, t) = \widetilde{A_\psi}\nabla\psi + \widetilde{A_\theta}\nabla\theta + \widetilde{A_\zeta}\nabla\zeta \tag{3.4}$$

and the perturbed electrostatic potential $\widetilde{\phi}(\boldsymbol{x}, t)$, we can write the energetic particle Lagrangian,

$$\mathcal{L} = \left(\rho_\parallel I + \chi + \widetilde{A_\theta}\right)\dot{\theta} + \left(\rho_\parallel g - \psi + \widetilde{A_\zeta}\right)\dot{\zeta} + \mu\dot{\xi} - \mathcal{H} + \left(\delta\rho_\parallel + \widetilde{A_\zeta}\right)\dot{\zeta} \tag{3.5}$$

with $\rho_\parallel = v_\parallel/\omega_{\mathrm{ci}}$ the 'parallel gyroradius', χ the toroidal flux, and g the metric tensor.

Therefore the canonical coordinates can be identified,

$$P_\theta = \rho_\parallel I + \chi + \widetilde{A_\theta} \tag{3.6}$$

$$P_\zeta = \rho_\parallel g - \psi + \widetilde{A_\zeta} \tag{3.7}$$

$$P_\xi = \mu \tag{3.8}$$

The particle guiding centre equations can therefore be derived, giving,

$$\dot{\theta} = \frac{1}{D}\left[\rho_\parallel B^2(1 - \rho_\parallel g' - \widetilde{A'_\zeta}) + g\left\{(\rho_\parallel^2 B + \mu)B' + \widetilde{\phi}'\right\}\right] \tag{3.9}$$

$$\dot{\zeta} = \frac{1}{D}\left[\rho_\parallel B^2(\rho_\parallel I' + q + \widetilde{A'_\theta}) - I\left\{(\rho_\parallel^2 B + \mu)B' + \widetilde{\phi}'\right\}\right] \tag{3.10}$$

$$\dot{\psi} = \frac{1}{D}\left[\left(I\frac{\partial\widetilde{A_\zeta}}{\partial\theta} - g\frac{\partial\widetilde{A_\theta}}{\partial\theta}\right)\dot{\theta} + \left(I\frac{\partial\widetilde{A_\zeta}}{\partial\zeta} - g\frac{\partial\widetilde{A_\theta}}{\partial\zeta}\right)\dot{\zeta} + g\dot{P}_\theta - I\dot{P}_\zeta\right] \tag{3.11}$$

$$\dot{\rho}_{\parallel} = \frac{1}{I}\left[\dot{P}_\theta - \frac{\partial \widetilde{A_\theta}}{\partial \theta}\dot{\theta} - \frac{\partial \widetilde{A_\theta}}{\partial \zeta}\dot{\zeta} - \frac{\partial \widetilde{A_\theta}}{\partial t} - \left(q + \frac{\partial \widetilde{A_\theta}}{\partial \psi} + \rho_{\parallel} I'\right)\dot{\psi}\right] \tag{3.12}$$

with $D = \rho_{\parallel}[gI' - Ig'] + I + qg - I\widetilde{A'_\zeta} + g\widetilde{A'_\theta}$, having chosen to use equations for $\dot{\psi}$ and $\dot{\rho}_{\parallel}$ instead of $\dot{P}_\theta$ and $\dot{P}_\zeta$. These four equations are then solved for each marker.

The equations for the markers must also be coupled with equations for the fields, calculated by evaluting the equations for the time derivatives. To study low-β shear Alfvén waves, we can choose the ordering to neglect $\delta B_{\parallel}$ and use the MHD assumption $E_{\parallel} = 0$. Later work [78] has dropped the restriction of $E_{\parallel} = 0$ when taking eigenfunctions from a gyrokinetic model, although that is not used in this work. These restrictions constrain $\widetilde{A_\perp}$, and $\widetilde{\boldsymbol{A}}$ can be rewritten as

$$\widetilde{\boldsymbol{A}} = \widetilde{\alpha}(\boldsymbol{x}, t)\boldsymbol{B} \tag{3.13}$$

This allows the field perturbations to be described using only a single scalar field, using,

$$E_{\parallel} = -\nabla_{\parallel}\widetilde{\phi} - \frac{\partial}{\partial t}(\widetilde{\alpha} B_0) = 0 \tag{3.14}$$

For each Alfvén mode, k, there remain two degrees of freedom, the slowly varying real amplitude (A_k) and mode phase (σ_k). The spatial structure can be decomposed poloidally into a series of poloidal Fourier harmonics,

$$\widetilde{\phi}_k = A_k(t)\mathrm{e}^{-i\sigma_k(t)} \sum_m \widetilde{\phi}_{km}(\psi)\mathrm{e}^{i(n_k\zeta - m\theta - \omega_k t)} \tag{3.15}$$

where $A_k(t)\mathrm{e}^{-i\sigma_k(t)}$ evolves slowly compared to the wave frequency.

Writing the interaction Lagrangian for a series of particles and some perturbed fields,

$$\mathcal{L}_{int} = \sum_{j=1}^{N_p}(\widetilde{\boldsymbol{A}}_j \cdot \boldsymbol{v}_j - \widetilde{\phi}_j) \tag{3.16}$$

By reformulating the perturbed amplitude from amplitude and phase to real and imaginary parts,

$$A_k(t)\mathrm{e}^{-i\sigma_k(t)} = \mathcal{X}(t) - i\mathcal{Y}(t) \tag{3.17}$$

The perturbed electrostatic potential at the position of the jth particle,

$$\widetilde{\phi}_j = \sum_{k=1}^{N_w}\sum_m [\mathcal{X}_k(t)C_{jkm} + \mathcal{Y}_k(t)S_{jkm}] \tag{3.18}$$

with $C_{jkm} \equiv \Re(\widetilde{\phi}_{km}(\psi_j)\mathrm{e}^{i\Theta_{jkm}})$ and $S_{jkm} \equiv \Im(\widetilde{\phi}_{km}(\psi_j)\mathrm{e}^{i\Theta_{jkm}})$, $\Theta_{jkm} \equiv n_k\zeta_j - m\theta_j - \omega_k t$.

Therefore we may write the integration Lagrangian,

$$\mathcal{L}_{int} = \sum_{j=1}^{N_p}\sum_{k=1}^{N_w}\frac{1}{\omega_k}\sum_m(k_{\|m}v_{\|j} - \omega_k) \tag{3.19}$$

Now, using also the wave Lagrangian,

$$\mathcal{L}_w = \sum_{k=1}^{N_w}\frac{E_k}{\omega_k}[A_k^2\dot{\sigma}_k] \tag{3.20}$$

with $E_k = 1/(2\mu_0)\int_V \left|\nabla_\perp\widetilde{\phi}_k\right|^2/v_A^2\mathrm{d}^3x$ we can vary $\mathcal{L}_{int} + \mathcal{L}_w$ with respect to $\mathcal{X}_k$ and $\mathcal{Y}_k$, we can come to the wave equations,

$$\dot{\mathcal{X}}_k = \frac{1}{2E_k}\sum_{j=1}^{N_p}\sum_m(k_{\|m}v_{\|j} - \omega_k)S_{jkm} \tag{3.21}$$

$$\dot{\mathcal{Y}}_k = -\frac{1}{2E_k}\sum_{j=1}^{N_p}\sum_m(k_{\|m}v_{\|j} - \omega_k)C_{jkm} \tag{3.22}$$

where we can clearly see the importance of the wave-particle resonances $(k_{\|m}v_{\|j} - \omega_k)$ for the evolution of the mode amplitude.

HAGIS then evolves the mode amplitudes using equations 3.21 and 3.22 together with the marker information, and the markers are evolved nonlinear using equations 3.9–3.12. For numerical efficiency, HAGIS employs an optional δf method, in which the distribution function is split into a constant background F_0 and a fluctuating part δf. This reduces numerical noise, as only the fluctuating part (typically small) needs to be discretized with markers. This entails some minor modifications to the equations solved. The full-f model is recovered by setting $F_0 = 0$, which gives $\delta f = f$.

We note also that the Lagrangian structure easily provides a scheme with energetic conservation.

3.2 Linear gyrokinetic eigenvalue solver LIGKA

The linear gyrokinetic eigenvalue solver LIGKA solves the linearized gyrokinetic equations to find the eigenvalues (frequency and damping) and eigenfunctions (mode structure) of the system. The model [79, 80] is concisely presented in reference [81], which

we briefly follow here.

We start with the linear gyrokinetic quasi-neutrality condition (QNE),

$$0 = \sum_s e_s \int \mathrm{d}^2 v \left\{ J_0 \delta f \right\}_s + m_\mathrm{i} \nabla_\perp \cdot \frac{n_\mathrm{i} \nabla_\perp \delta\phi}{B^2} + \frac{3 P_{\mathrm{i}\perp}}{4 B^2 \Omega_\mathrm{i}^2} \nabla_\perp^4 \delta\phi \tag{3.23}$$

and the linear gyrokinetic momentum equation (GKM),

$$\begin{aligned} -\frac{\partial}{\partial t}\left[\nabla \cdot \frac{\nabla_\perp \delta\phi}{v_\mathrm{A}^2}\right] + \boldsymbol{B} \cdot \nabla \frac{\nabla \times (\nabla \times \delta A_\parallel \boldsymbol{b})}{B} + (\boldsymbol{b} \times \nabla \delta A_\parallel) \cdot \nabla \frac{\mu_0 j_\parallel}{B} \\ = -\sum_s \mu_0 \int \mathrm{d}^2 v \, e_s \left\{ \boldsymbol{v}_\mathrm{d} \cdot \nabla J_0 \delta f \right\}_s + \sum_s \left[\boldsymbol{b} \times \nabla \left(\frac{\beta_{s\perp}}{2\Omega_s} \right) \right] \cdot \nabla \nabla_\perp^2 \delta\phi \\ + \sum_s \frac{3\beta_{s\perp}}{8\Omega_s^2} \nabla_\perp^4 \frac{\partial \phi}{\partial t} + \boldsymbol{B} \cdot \nabla \frac{1}{B} \sum_s \frac{\beta_s}{4} \nabla_\perp^2 \delta A_\parallel \end{aligned} \tag{3.24}$$

By splitting off the adiabatic part of the distribution function, and separating the equilibrium part and a perturbed part

$$\delta f_s = \delta h_s + \frac{\partial f_{eq,s}}{\partial E} e_s J_0 (\rho k_\perp) \left[\delta\phi - \delta\psi - \frac{v_\parallel k_\parallel}{\omega} \delta\psi \right] + \frac{\nabla f_{eq,s}}{i\omega B} \cdot (\boldsymbol{b} \times \nabla) J_0 \delta\psi \tag{3.25}$$

We can then write the linear gyrokinetic equation (GKE),

$$\frac{\partial \delta h}{\partial t} + (v_\parallel \boldsymbol{b} + \boldsymbol{v}_\mathrm{d}) \cdot \nabla \delta h = \left[\frac{\boldsymbol{b} \times \nabla f_\mathrm{eq}}{eB} \cdot \nabla - \frac{\partial f_\mathrm{eq}}{\partial E} \frac{\partial}{\partial t} \right] J_0 \left[\delta\phi - \left(1 - \frac{\boldsymbol{v}_\mathrm{d} \cdot \nabla}{i\omega} \right) \delta\psi \right]. \tag{3.26}$$

with $\boldsymbol{v}_d = -\boldsymbol{b}/(eB) \times \left(m v_\parallel^2 (\boldsymbol{b} \cdot \nabla) \boldsymbol{b} + \mu \nabla B \right)$ the drift velocity, ϕ the electrostatic potential, ψ the electromagnetic superpotential ($\partial_t \delta A_\parallel(\boldsymbol{x}) = -\boldsymbol{b} \cdot \nabla \delta\psi(\boldsymbol{x})$), and $J_0 = J_0(\rho k_\perp)$ the gyroaverage operator.

By solving these equations, and integrating over the particle orbits as necessary, it is possible to solve the eigenvalue problem for linear global gyrokinetic treatment of modes with frequencies spanning from zero-frequency MHD-type global modes such as kink instabilities to all types of shear Alfvén waves (SAW), but excluding high frequency compressional Alfvén eigenmodes.

The system of equations are discretized radially with finite elements, and decomposed via Fourier representation into a toroidal mode number n and a summation of poloidal harmonics m.

Here, as in reference [81], we primarily use the "analytical version" of LIGKA

unless otherwise stated, namely approximating the integration along the orbits of particles by analytical expressions. This leads to the fast-circulating approximation, namely neglecting effects of the trapped particles. In the case of the electron Landau damping, for example, where the damping comes primarily from the particles close to the trapped-passing boundary, this can lead to underestimations away from the rational surfaces and overestimations close to the rational surfaces. This analytical version has demonstrated that it is able to reproduce the results from the analytical dispersion relation of reference [82] in reference [83] without finite Larmor radius effects, and including these effects it has reproduced the dispersion relation of reference [84] in reference [80].

3.2.1 Limits

It is possible to solve the equations in the limit of reduced MHD where the non-adiabatic part of the distribution function is neglected ($h = 0$), the electrostatic and electromagnetic potentials are set equal ($\delta\phi = \delta\psi$), $E_{\|} = 0$, and the gyroaveraging is neglected. This allows us to compare the effects of a kinetic treatment with an MHD treatment.

Whereas the global calculations provide the radial structure of each of the poloidal harmonics, one can also run the code in the local limit. Here, the system is solved only at a single radial position. This is very fast, and gives a relatively accurate estimate for the mode drive (if included), and for the Landau damping. The continuum damping is neglected, which we know becomes weak for modes localized in the gap. Whilst one can also apply a model which allows the radiative damping to be estimated [80], this is not included in this work.

3.2.2 Workflows with HAGIS and LIGKA

Before we can perform simulations with HAGIS, we need to obtain a mode structure, frequency, and damping, all of which we take from LIGKA. As previously mentioned, there are different models available within LIGKA. We therefore outline here the three workflows which we shall use with HAGIS and LIGKA. In all cases, as the interaction with EPs is included in HAGIS, we run LIGKA without any EPs present. The two global workflows are sketched in figure 3.1.

Global kinetic (antenna) The highest fidelity mode structure calculation that we use from LIGKA for use in HAGIS is calculated using the global gyrokinetic model of LIGKA, the main working model of the code. In this case, we run automatically the following chain of calculations with LIGKA, the output of each step acting as a

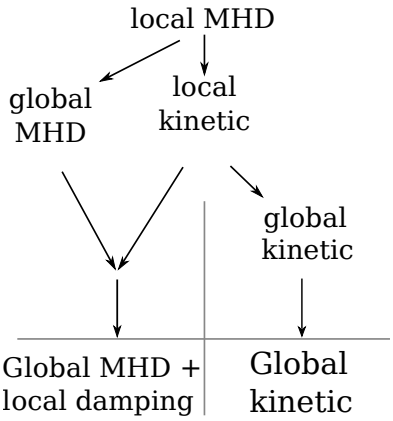

Figure 3.1: Global workflows used for providing mode structure, frequency, and damping from LIGKA for use in HAGIS.

starting point for the next. First, we run the local ideal MHD solver, which returns the crossing point of the continuum in the ideal limit. Next, we run the local kinetic solver, which returns the complex frequency of the tip of the continuum. Finally, we run the global kinetic solver, which returns the global eigenfunction and the complex mode frequency, obtained by first scanning the frequencies across the gap to find any peaks in the response, before refining the scan around any peaks found. In cases where multiple eigenmodes are found, we sort the modes by damping, and select the mode with the weakest damping. Note that as the method for the kinetic solver is based on the response to a frequency sweep, this method is also referred to as the antenna method.

Global MHD A somewhat reduced model is to run LIGKA in its MHD limit. In this case, we first run the local ideal MHD solver, before running the global MHD eigenvalue solver. This is solved by means of a full matrix inversion, which returns all eigenvalues and eigenfunctions of the matrix. The eigenvalues are then sorted according to frequency and we select the most appropriate candidate, the mode with the closest frequency with the correct mode properties. The eigenvalue (the real frequency) and the eigenfunction are then returned. As the MHD model does not include any damping, we therefore also run the local kinetic solver, which we take as the imaginary part of our frequency. As these steps are relatively fast, we have modified HAGIS to directly make these calls to the library version of LIGKA, creating a directly coupled interface. If the mode structure or damping calculations have already been performed, then we use the cached versions to save computational resources.

Local The simplest model is to run LIGKA only in its local limits. We first run the local ideal MHD solver, before running the local kinetic solver. Using an analytical estimate for the mode structure width [85], we then use as an eigenfunction two poloidal harmonics with Gaussian radial structure. Again, this model is implemented in a directly coupled interface.

3.3 Nonlinear initial value gyrokinetic code ORB5

The nonlinear global electromagnetic gyrokinetic code ORB5 [46, 86, 87] implements the most complete physical model outlined. It uses the Monte Carlo Lagrangian Particle-in-Cell (or PIC) method for evolving the distribution function, sampled using markers, and the fields (electrostatic and electromagnetic potentials) are solved on a grid using finite element representation.

The model of the code, constructed from a nonlinear global electromagnetic gyrokinetic Lagrangian [88], is covered in detail in references [46], [54], and [89], however we here outline some of the key parts, focussing on the details most relevant for this work, following those references.

The distribution function for each species s is first split into a constant background part, F_{0s} and a varying part, δf_s. By representing only the time varying part of the distribution with Monte Carlo markers, the quality of the signal with respect to numerical noise is greatly improved.

$$\delta f_s(\boldsymbol{R}, v_\parallel, \mu, t) = \sum_{i=1}^{N_p} w_{s,i}(t)\delta(\boldsymbol{R} - \boldsymbol{R}_i)\delta(v_\parallel - v_{\parallel,i})\delta(\mu - \mu_i) \tag{3.27}$$

The perturbed part of the distribution function is then evolved according to the gyrokinetic Vlasov equation

$$\frac{\partial \delta f_\mathrm{s}}{\partial t} + \dot{\boldsymbol{R}} \cdot \left.\frac{\partial \delta f_\mathrm{s}}{\partial \boldsymbol{R}}\right|_{v_\parallel} + \dot{v}_\parallel \frac{\partial \delta f_\mathrm{s}}{\partial v_\parallel} = -\dot{\boldsymbol{R}} \cdot \left.\frac{\partial F_{0s}}{\partial \boldsymbol{R}}\right|_{\epsilon} - \dot{\epsilon}\frac{\partial F_{0s}}{\partial \epsilon} \tag{3.28}$$

according to the equations for the particle equations of motion

$$\dot{\boldsymbol{R}}^{(0)} = v_\parallel \boldsymbol{b} - v_\parallel^2 \frac{c m_\mathrm{s}}{q B_\parallel^*}\boldsymbol{G} + \mu \frac{B c m_\mathrm{s}}{q B_\parallel^*}\boldsymbol{b} \times \frac{\nabla B}{B} \tag{3.29}$$

$$\dot{\boldsymbol{R}}^{(1)} = \frac{\boldsymbol{b}}{B_\parallel^*} \times \nabla\left\langle \phi - v_\parallel A_\parallel^\mathrm{s} - v_\parallel A_\parallel^\mathrm{h} \right\rangle - \frac{q_\mathrm{s}}{m_\mathrm{s}} \langle A_\parallel^\mathrm{h} \rangle \boldsymbol{b}^* \tag{3.30}$$

$$\dot{v}_{\parallel}^{(0)} = \mu B \nabla \cdot \boldsymbol{b} + \mu v_{\parallel} \frac{c m_{\mathrm{s}}}{q_{\mathrm{s}} B_{\parallel}^{*}} \boldsymbol{G} \cdot \nabla B \tag{3.31}$$

$$\dot{v}_{\parallel}^{(1)} = -\frac{q_{\mathrm{s}}}{m_{\mathrm{s}}} \left[\boldsymbol{b}^{*} \cdot \nabla \left\langle \phi - v_{\parallel} A_{\parallel}^{(h)} \right\rangle + \frac{\partial}{\partial t} \langle A_{\parallel}^{(s)} \rangle \right] - \mu \frac{\boldsymbol{b} \times \nabla B}{B_{\parallel}^{*}} \cdot \nabla \langle A_{\parallel}^{(s)} \rangle \tag{3.32}$$

$$\dot{\mu} = 0 \tag{3.33}$$

$$\dot{\epsilon}^{(0)} = 0 \tag{3.34}$$

$$\dot{\epsilon}^{(1)} = v_{\parallel} \dot{v}_{\parallel}^{(1)} + \mu \nabla B \cdot \dot{\boldsymbol{R}}^{(1)} \tag{3.35}$$

with

$$\boldsymbol{B} = \nabla \times \boldsymbol{A} \tag{3.36}$$

$$\boldsymbol{b} = \boldsymbol{B} / B \tag{3.37}$$

$$B_{\parallel}^{*} = \boldsymbol{b} \times \boldsymbol{A}^{*} \tag{3.38}$$

$$\boldsymbol{A}^{*} = \boldsymbol{A} + \left(\frac{m_{\mathrm{s}} c}{q_{\mathrm{s}}} v_{\parallel} + \langle A_{\parallel}^{(s)} \rangle \right) \boldsymbol{b} \tag{3.39}$$

$$\boldsymbol{b}^{*} = \frac{\nabla \times \boldsymbol{A}^{*}}{B_{\parallel}^{*}} = \boldsymbol{b} - \left(\frac{c m_{\mathrm{s}}}{q B_{\parallel}^{*}} v_{\parallel} + \langle A_{\parallel}^{(s)} \rangle \right) \boldsymbol{G} \tag{3.40}$$

$$\boldsymbol{G} = \boldsymbol{b} \times (\boldsymbol{b} \times (\nabla \times \boldsymbol{b})) \tag{3.41}$$

and with the gyroaveraged potential $\langle \phi \rangle = \oint \phi(\boldsymbol{R} + \boldsymbol{\rho}) \mathrm{d}\alpha / (2\pi)$ with ρ the gyroradius of the particle and α the gyro-phase.

These equations are coupled to the field equations, the gyrokinetic quasineutrality equation and the parallel Ampère's law. Firstly the gyrokinetic quasineutrality equation,

$$-\nabla \cdot \left[\left(\sum_{s} \frac{q_{\mathrm{s}}^{2} n_{\mathrm{s}}}{T_{\mathrm{s}}} \rho_{\mathrm{s}}^{2} \right) \nabla_{\perp} \phi \right] = \sum_{s} q_{\mathrm{s}} n_{1\mathrm{s}} \tag{3.42}$$

The perturbed magnetic potential $A_{\parallel}$ is split in to the symplectic part $A_{\parallel}^{(s)}$, which is

found from

$$\frac{\partial}{\partial t} A_{\parallel}^{(s)} + \boldsymbol{b} \cdot \nabla \phi = 0 \tag{3.43}$$

and the Hamiltonian part $A_{\parallel}^{(h)}$, solved using the mixed-varible parallel Ampère's law,

$$\left(\sum_s \frac{\beta_\mathrm{s}}{\rho_\mathrm{s}^2} - \nabla_\perp^2 \right) A_{\parallel}^{(h)} = \mu_0 \sum_s j_{\parallel 1s} + \nabla_\perp^2 A_{\parallel}^{(s)} \tag{3.44}$$

with $n_{1\mathrm{s}} = \int \mathrm{d}^6 Z \delta f_\mathrm{s} \delta(\boldsymbol{R} + \boldsymbol{\rho} - \boldsymbol{x})$ the perturbed gyrocenter density, $\rho_\mathrm{s} = \sqrt{m_\mathrm{s} T_\mathrm{s}}/(q_\mathrm{s} B)$ the thermal gyroradius, q_s the particle charge, $\mathrm{d}^6 Z = B_{\parallel}^* \mathrm{d}\boldsymbol{R} \mathrm{d}v_{\parallel} \mathrm{d}\mu \mathrm{d}\alpha$ the phase space volume, $j_{1\mathrm{s}} = \int \mathrm{d}^6 Z v_{\parallel} \delta f_\mathrm{s} \delta(\boldsymbol{R} + \boldsymbol{\rho} - \boldsymbol{x})$ the perturbed parallel gyrocenter current. Note that due to the scale separation between ion and electron Larmor radii, electrons are treated differently from ion species, and terms with ρ_e are neglected.

The particular method of using the mixed-variables formulation, and then applying a pullback procedure to the particle weights (for linear simulations), is detailed in reference [54]. The idea is that by splitting $A_{\parallel}$ into the symplectic part and the Hamiltionian, and using the fact that we can make a good prediction for the evolution of the symplectic part using Ohm's law (equation 3.43), that we can significantly improve the behaviour of the electromagnetic part of the field solution, thus allowing us to take larger steps at the same error, or much smaller error for the same step size. At the end of the timestep, the total of the $A_{\parallel}$ is moved to the symplectic part, and the Hamiltonian part is set to zero. Then, the particle weights are modified according to the "pullback" transformation such that the distribution function is not modified by this change of representation.

$$\delta f_{s(new)}^{(m)} = \delta f_s^{(s)} = \delta f_{s(old)}^{(m)} + \frac{q_s \langle A_{\parallel(old)}^{(h)} \rangle}{m_s} \frac{\partial F_{0s}}{\partial v_{\parallel}} \tag{3.45}$$

where the superscripts (m) and (s) on the distribution function refer to the mixed and symplectic representations respectively.

For nonlinear simulations, the transformation is slightly different, thus we apply

$$v_{\parallel(new)}^{(m)} = v_{\parallel}^{(s)} = v_{\parallel(old)}^{(m)} - \frac{q_s}{m_s} \langle A_{\parallel}^{(h)} \rangle \tag{3.46}$$

$$f_{1s(new)}^{(m)} \left(v_{\parallel(new)}^{(m)} \right) = f_{1s(old)}^{(m)} \left(v_{\parallel(old)}^{(m)} \right) \tag{3.47}$$

instead of equation 3.45.

A set of straight field line coordinates are used, with radial coordinate $s = \sqrt{\psi/\psi_\mathrm{edge}}$

(where ψ is the poloidal flux), toroidal angle φ, and poloidal angle

$$\chi = \frac{1}{q(s)} \int_0^\theta \frac{\boldsymbol{B} \cdot \nabla \varphi}{\boldsymbol{B} \cdot \nabla \theta'} \mathrm{d}\theta'$$

(where θ is the geometric poloidal angle).

Linear and nonlinear models The model above is a nonlinear model. However, the terms labelled with superscript '(1)' are nonlinear terms which depend on the perturbed fields. Therefore, by neglecting these terms, we allow a species to follow only the unperturbed particle trajectories.

As we discussed before, the wave-particle (WP) nonlinearities are retained by allowing the energetic particles to perturb their trajectories.

One can draw a distinction between the roles of different nonlinearities present in the system. The nonlinear interaction between the wave and the energetic particles present in the system is referred to as the wave-particles (WP) nonlinearity [90]. To capture this effect, it is sufficient to retain the nonlinearity only on the energetic particle species. This is the same nonlinearity as is present in HAGIS. The other nonlinearity present in the system is associated with the nonlinearities of the bulk plasma species, and corresponds to the nonlinear interaction of waves with each other (in particular via three-wave coupling) and is therefore known as the wave-wave (WW) nonlinearity. The most important effect of WW nonlinearities is to couple different modes together, creating a channel to drive (stable) low n modes, especially the zonal $n = 0$ mode. We note here the difficulties associated with low mode number global electromagnetic simulations of high-β plasmas, which prevent us from retaining low ($n \lesssim 5$) toroidal mode numbers in the ITER studies. We note some progress on this subject has recently been made [91], but this issue is still outstanding for the parameters involved in this study. Inclusion of these effects will likely have the effect of reducing the saturation amplitude of the modes.

Given that performing simulations with WP nonlinearities enabled are not significantly more difficult in ORB5 than linear simulations, it is of course also reasonable to analyse the linear behaviour of a system by using the nonlinear model, during the linear phase where the nonlinear terms are small. The results presented later for the ITER cases with shaped profiles will be based on analysis of the linear phase of simulations retaining the WP nonlinearity.

Discretization As previously mentioned, the perturbed distribution function ($\delta f_s = f_s - F_{0s}$) is discretized with markers. The fields are discretized with B-Splines (in this work cubic). Once the density (and current) deposition from the markers to the grid

has been performed, Fourier transforms in toroidal and poloidal angles are applied, and the fields are filtered according to a radially dependent filter for each toroidal mode number n, the poloidal modes kept, $m_n(r) \in \lfloor nq(r) \rfloor \pm \Delta m$, with Δm typically 5.

Parallelization ORB5 is a massively parallel code, and has long been parallelized with MPI using domain decomposition and domain cloning. Recent work [92] has extended this to also include hybrid MPI/OpenMP or MPI/OpenACC parallelization. Due to the large radial and poloidal resolution used in this work, this extension is immensely useful, and allows us to obtain good performance when scaling the code beyond the domain decomposition, which is limited by the number of toroidal points. Typical simulations in this work use OpenMP parallelization with the number of threads of order 2 – 6, depending on the case and the configuration of the computer where the simulation is performed. For presented simulations with many toroidal mode numbers, where we use very large numbers of markers, the parallelization is increased up to cases with 24 threads.

Normalization Various quantities, both in the input and output of the code are subject to normalizations. Therefore these must be calculated when setting up simulations of a particular scenario, and when interpreting simulations results. These quantities are based on the ion mass m_i, ion charge $q_\mathrm{i} = eZ_\mathrm{i}$, magnetic field on axis B_0, the electron temperature at a reference radius (in this work always chosen to be on axis) $T_\mathrm{e}(0)$, and the average electron density $\overline{n_\mathrm{e}}$.

Time scales are normalized to the ion cyclotron frequency, $\omega_\mathrm{ci} = q_\mathrm{i} B_0 / m_\mathrm{i} c$. When measuring frequencies of Alfvénic modes simulated with ORB5, we shall convert these frequencies to units of the Alfvén frequency, $\omega_\mathrm{A} = v_\mathrm{A}/R$, noting that for ITER parameters, the factor $\omega_\mathrm{ci}/\omega_\mathrm{A} \lesssim 200$. Length scales are normalized to the ion sound Larmor radius, $\rho_\mathrm{s} = c_\mathrm{s}/\omega_\mathrm{ci}$, with $c_\mathrm{s} = \sqrt{eT_\mathrm{e}(0)/m_\mathrm{i}}$, noting that for ITER parameters the horizontal equilibrium dimension $2a \approx 850\rho_\mathrm{s}$.

Diagnostics The majority of the analyses presented in this work, will be based on the electrostatic and electromagnetic potentials. During this work, the diagnostic output of the code has been improved, and now in addition to poloidal and toroidal slices, the fields are also output in their full Fourier representation. Because the field solvers apply Fourier filtering, this reduces significantly the amount of data required to completely represent the three dimensional fields, allowing us to perform diagnostics on the full field data, necessary to be able to simultaneously resolve toroidal and poloidal mode numbers. For many diagnostics, we prefer to keep the analysis in the Fourier representation, for example when plotting poloidal harmonic amplitudes. The real space signal is

reconstructed if we, for example, wish to plot the potential on a poloidal cross-section. Frequency analysis (unless otherwise stated) is performed on the electrostatic potential at the outer mid-plane (poloidal angle $\chi = 0$).

An interesting diagnostic development is the ability to calculate the species and phase space resolved power exchange [93]. However, the meaning of such a method depends on the coordinates used, and therefore this needs to be reformulated for use in electromagnetic simulations using the mixed-variables fomulation. Progress on this has been made [89], but the extension to real or phase space resolved data is outstanding.

3.4 Discussion of numerical tools

In this chapter, we have introduced a hierarchy of models for performing linear and nonlinear simulations of Alfvénic physics. The most complete numerical model available is the nonlinear, non-perturbative global electromagnetic gyrokinetic code ORB5, which has, however, prior to this work never been applied to studies of Alfvén eigenmodes in scenarios with parameters as challenging as those in this work. Simulations using such a model are difficult, and a methodical validation procedure is therefore performed as we begin to run linear simulations for the ITER scenario.

On the other hand, linear results from the gyrokinetic eigenvalue code LIGKA are relatively robust, and the code has previously been extensively validated through benchmarks and experimental comparisons. However, its applicability is restricted to linear physics.

The nonlinear perturbative HAGIS model, coupled with eigenfunctions from LIGKA is able to evaluate linear stability calculations (in the drift-kinetic limit), and assess mode saturation due to wave-particle nonlinear saturation. In this case, its applicability is limited by the perturbative assumption, and, as the modes must be prescribed *a priori*, one must already predict which modes will be relevant. As an example, as we shall see later, it is not obvious from the linear studies, in chapters 4 and 5 that one should retain odd-parity TAE modes. However, as we find in chapter 6, these modes may nonlinearly become significant.

Numerically, linear, single-mode simulations with ORB5 can cost $\sim 10^4$ core-hours, with the most expensive multi-mode nonlinear simulations costing $\sim 10^6$ core-hours, whereas linear hybrid simulations can be performed in a few minutes, and nonlinear simulations can run in a few hours. Also multi-mode hybrid simulations can become very numerically demanding, with an estimate in the scaling with respect to number of modes $\sim n^2_{\text{modes}}$.

4 Perturbative hybrid-kinetic modelling of TAEs in ITER

4.1 Introduction

In this chapter, we aim to address the issue of the TAEs in the ITER scenario, which we do using the fast and automated workflows previously introduced. We revisit a scenario which has been previously studied [20], although we perform a systematic study with many modes, spanning a broad range of toroidal mode numbers, TAE gaps, and we additionally include also the odd-parity TAEs, and look at also the subdominant modes. In addition to the regular scenario, we also include an additional scenario, based on a flat q profile, taking a slightly different approach to investigating the profile sensitivity as was performed in reference [94]. We also appreciate the issue of understanding the limitations and sensitivity of any results, which is why we include a systematic scan of the effects of the number of poloidal harmonics. Finally, we also perform tests of a purely local workflow, which we might expect to work for cases with very localized mode structures, such as those seen for large toroidal mode numbers.

4.2 Linear results

Taking the standard ITER scenarios, as well as an additional flat profile scenario, we investigate the linear properties of the TAE modes. We perform a systematic study of multiple TAE gaps, scanning the toroidal mode number n, and the mode parity.

For the standard scenario, as q is close to 1, we take the 'zeroth' gap, where $m = (n, n+1)$, as well as the next ('first') gap, $m = (n+1, n+2)$. For the flat q scenario, we take the 'zeroth' and 'first' gaps, and as well as the previous ('minus-oneth') gap, $m = (n-1, n)$, due to the lower values of the q-profile.

Figure 4.1 (left) shows the mode drive, frequency, damping, and total growth rate for the two gaps using both workflows for the standard scenario. In the case of the global kinetic workflow, this is performed for both even and odd mode parity, and shown against the global MHD workflow (even parity). One can observe a good agreement between the workflows in the drive for large n. We also note a good agreement for

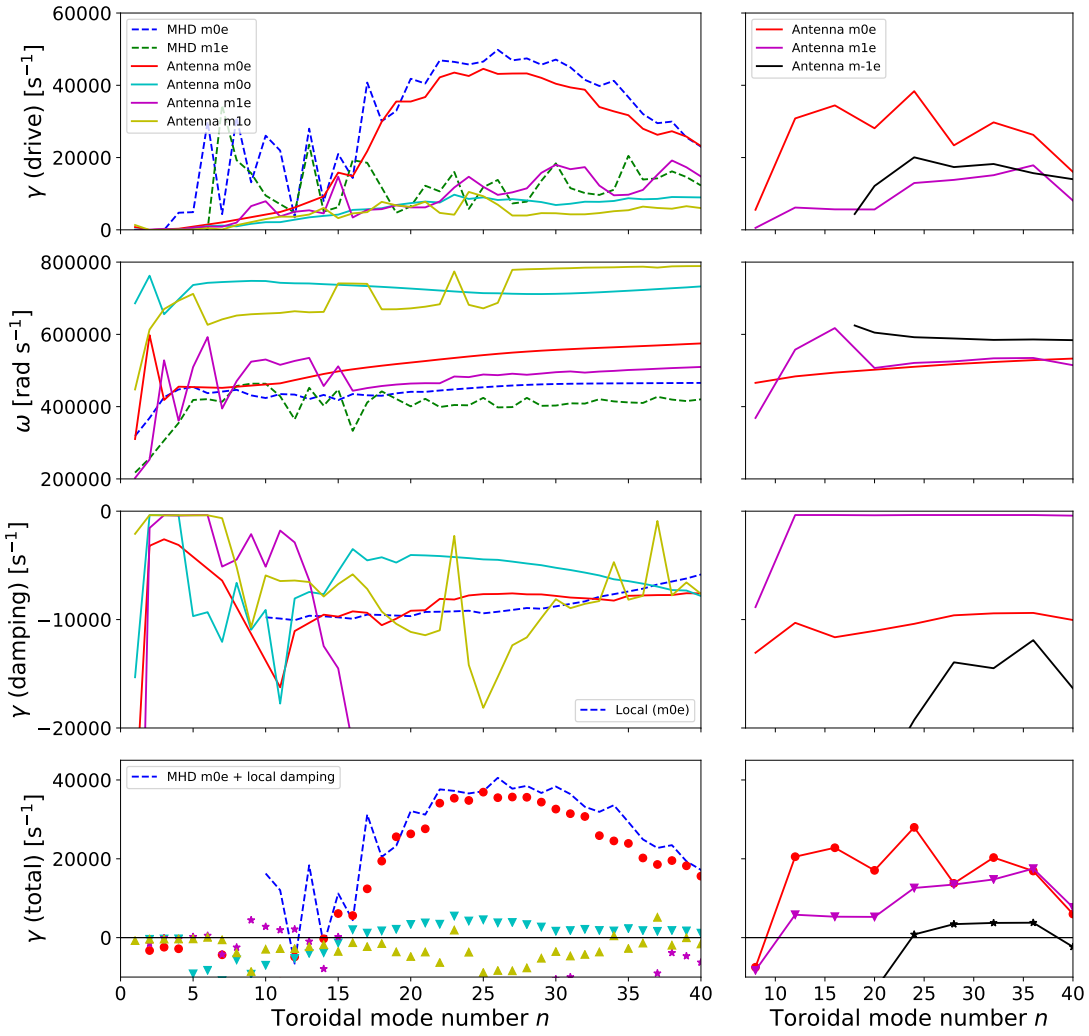

Figure 4.1: Linear TAE properties for (left) the standard ITER equilibrium and (right) the flat q ITER equilibrium. The dashed lines in the left figure correspond to the global MHD workflow and the solid lines (and points) refer to the global kinetic (antenna) workflow. In the standard case (left), the "zeroth" gap ($m = (n, n+1)$) and the "first" gap ($m = (n+1, n+2)$) are presented (labelled 'm0' and 'm1' respectively). In the flat q case (right) the "zeroth", "first", and "minus-oneth" gaps ($m = (n-1, n)$, labelled 'm-1') gaps are presented. In both cases, the suffices 'e' and 'o' refer to the even or odd parity of the TAE mode.

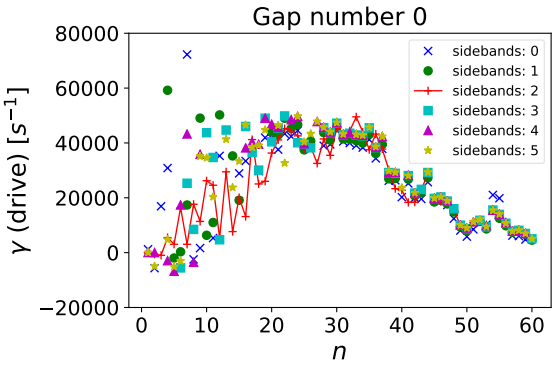

Figure 4.2: The TAE drive measured using the MHD workflow for varying numbers of poloidal sidebands included in the calculation. The number of sidebands (sb) indicated in the legend are both above and below the main pair of poloidal mode numbers, $n - sb \leq m \leq n + sb + 1$.

$n \gtrsim 15$ in the damping rates between that from the local kinetic calculation, as used in the MHD workflow, and the global kinetic eigenvalue.

Overall, we note that the even-parity 'zeroth' gap exhibits the largest total growth rates, although for lower mode numbers $n \approx 10$, we observe that the most unstable modes correspond to the (even) 'first' gap. We also see for medium mode numbers, ($n \approx 15 \rightarrow 40$) that the odd 'zeroth' gap modes are also weakly unstable, owing primarily to their weak damping rates.

Figure 4.1 (right) shows the linear properties of the TAEs for the flatter ITER scenario, showing the drive, frequency, damping, and total growth rate for three gaps ('zeroth', 'first', and 'minus-oneth'), all calculated using the global kinetic workflow.

For this scenario, we see that the 'zeroth' gap is most unstable, and along with the 'first' gap, is unstable over a wide range of mode numbers ($10 \lesssim n \lesssim 45$). However, for higher mode numbers ($25 \lesssim n \lesssim 40$), also the 'minus-oneth' branch is unstable.

For this scenario as with the standard scenario, the largest drive is seen in the 'zeroth' gap, however, when compared with the standard scenario, the most unstable toroidal mode numbers are reduced, with large drive seen already for $n = 12$. Also, for $n \gtrsim 15$, the damping rate is comparable for the 'zeroth' gap between the scenarios, however for the 'first' gap, the damping is much reduced compared with the standard scenario. We explain this due to a reduction in the ion Landau damping. As the modes are located at larger radial position, this corresponds to a position where the ion temperature is lower. As the ion Landau damping occurs primarily in the tail of the Maxwellian distribution, this term is exponentially sensitive to the ion temperature. This leads to the total growth rates being comparable for larger n between the 'zeroth' and 'first' gaps.

In the case of lower n ($\lesssim 15$), we see a large deviation in the drive between the

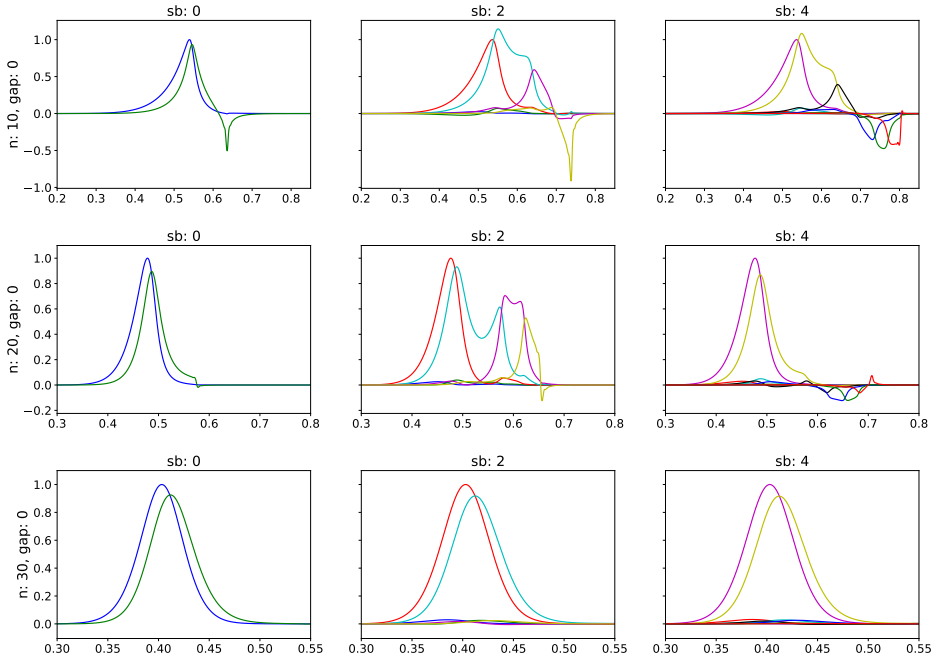

Figure 4.3: Radial TAE eigenfunctions from the MHD workflow calculated with different numbers of poloidal harmonics (left: 0 sidebands, 2 harmonics; middle: 2 sidebands, 6 harmonics; right: 4 sidebands, 10 harmonics) for toroidal mode numbers (top: 10, middle: 20, bottom: 30).

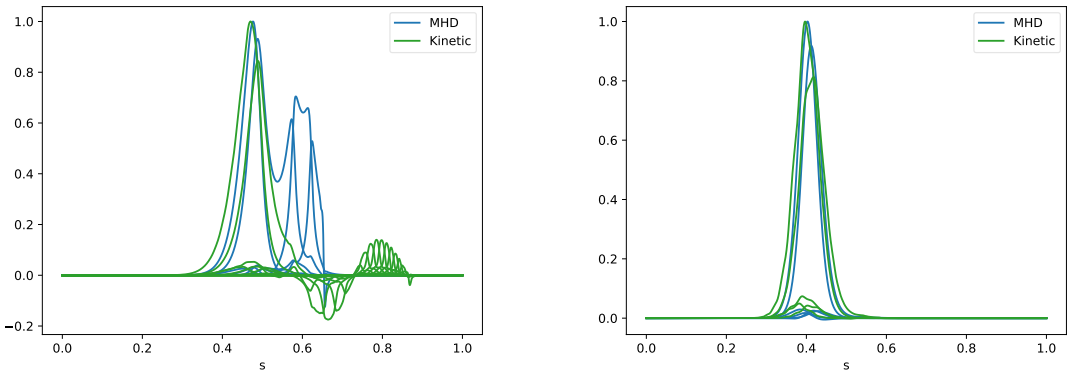

(a) **(b)**

Figure 4.4: TAE eigenfunctions calculated with global MHD workflows (blue) and global kinetic workflows (green) for (a) $n = 20$ and (b) $n = 30$. The MHD eigenfunctions were calculated with 2 sidebands, corresponding to the central column of figure 4.3.

MHD and kinetic workflows. In these cases, we experienced also difficulty in resolving these MHD workflow calculations. For lower n, it is necessary to increase the number of poloidal harmonics significantly for the solution to converge. Figure 4.2 shows the measured TAE drive as a function of the toroidal mode number n, for varying numbers of poloidal harmonics included in the global MHD workflow. We note that for $n \gtrsim 25$, this converges readily with respect to the number of sidebands. Due to the increased expense of solving the eigenvalue problem, we solve this up to 5 sidebands (12 poloidal harmonics), although we note that even in this case, strict convergence is not reached for all low mode numbers ($n \lesssim 20$). As the MHD workflow is not the main focus of this work, we do not pursue this further, but instead restrict our attention for this workflow to $n \gtrsim 20$.

Figure 4.3 shows the MHD eigenfunctions for different numbers of sidebands, for a range of n. It can be seen that for larger n (for example $n = 30$), the mode structure is narrow, without radially extended interaction of the TAE with the continuum. As a result, the mode structure is very simple, and can be accurately represented even with few harmonics retained. For lower n, the mode structures are more global, with features present in more harmonics, therefore a larger number of harmonics are needed to represent the mode structures accurately.

We shall find in the next chapter that these global, lower n modes are indeed confirmed by global gyrokinetic simulations to be weakly damped, and stress the importance of retaining their effects, in contrast to previous works which used local models [22] or neglected the global low n modes [95].

4.2.1 Influence of the energetic particle distribution function on the drive

In light of work to be shown in the next chapter, we take the opportunity to compare the drive between cases with the nominal distribution function for the alpha particles, an isotropic slowing down distribution with birth energy of 3.5 MeV, and a Maxwellian distribution with temperature of 900 keV [22, 96]. We choose two representative modes from the main gap, $n = 20$, and $n = 30$. The results, calculated in the absence of damping, including only the drive, are in table 4.1.

Table 4.1: Measured drive for TAEs driven by alpha particles with Maxwellian or slowing down distribution functions.

	Maxwellian	Slowing Down
n=20	$3.55{\cdot}10^4$ s^{-1}	$4.42{\cdot}10^4$ s^{-1}
n=30	$3.73{\cdot}10^4$ s^{-1}	$4.14{\cdot}10^4$ s^{-1}

This confirms that the error in the TAE drive is relative small, which justifies the simplification of the EP distribution function in the next chapter. This will also allow us to contextualize the growth rates measured with the Maxwellian distribution function in the next chapter, and can in principle be used to renormalize the observed growth rates for a comparison with those from a slowing down distribution.

4.3 Nonlinear saturation

Taking the standard ITER scenario from the above section, we now run longer simulations of the even parity zeroth gap (the gap with the largest linear growth rates), allowing the mode growth to saturate due to nonlinear wave-particle trapping, akin to what is seen in a one-dimensional bump-on-tail model [40]. Note that these simulations are still performed for single TAEs. We measure the saturation amplitude ($A_{\max}$) as the first maximum after the linear growth phase, measured in units of relative perturbed radial magnetic field ($\delta B_{\mathrm{r}} / |B|$). We perform this exercise for both MHD and kinetic workflows and the results are plotted in figure 4.5.

In figure 4.5a, we observe a peak in the saturation amplitudes at $n = 25$ for the global kinetic workflow (maximum amplitudes $\sim 3.5 \cdot 10^{-3}$), and a similar but broader peak for the global MHD workflow at somewhat lower amplitude ($\sim 2 \cdot 10^{-3}$). In figure 4.5b, we plot the saturation amplitude from figure 4.5a against the normalized linear growth rates from the previous section. Here we observe a scaling for both workflows which looks approximately quadratic, as one might expect from analytical theory [97]. Such analytical calculations, however, employ several simplifications, most noticeably assuming that the nature of the resonance is not changing as γ/ω changes, which is virtually fulfilled in the case of varying γ by varying the energetic particle density. On the other hand, we here change the linear drive and damping by changing the modes themselves, therefore both the mode numbers and the resonance properties are changing as well. For this reason, we do not expect any exact agreement with quadratic scaling. In figure 4.5c, we plot the coefficient of proportionality between the saturation amplitude and a quadratic scaling assumption as a function of the toroidal mode number.

The first conclusion to be drawn is that the coefficients differ for the two workflows. A possible reason for this would be the kinetic upshift of the frequencies as compared to the MHD results (seen in the second panel of figure 4.1), even though the measured drive and growth rates are broadly similar. An additional effect of the increased frequencies is a change in the population of energetic particles which are resonant with the mode. The next conclusion is that for the MHD workflow, the coefficient is close to constant (~ 0.3), implying that a quadratic scaling is indeed a reasonable approximation of the

relationship between the saturation amplitude and the linear mode properties. In the case of the the kinetic workflow, the coefficients are larger (~ 0.8) and show a weak positive correlation with n. Finally, we notice that other than at $n = 20$, where the reduced damping was barely sufficient to saturate the mode, the cases with significantly reduced damping in the MHD workflow show very similar scaling coefficients to the cases with nominal damping.

To summarize, we observe that, even though the quadratic scaling is not perfectly recovered in the case of the kinetic workflow, the deviation from it is smooth and therefore the entire range of saturation amplitudes can easily be predicted from the linear mode properties and from a small number of cases run nonlinearly and used to fit the quadratic scaling coefficient. These results can be used as input in the view of further development of reduced quasilinear EP transport models (e.g. in reference [98]).

4.4 Towards non-perturbative nonlinear saturation

Nonlinear results observed with non-perturbative hybrid models exhibit mode structure variation associated with frequency chirping [30], we shall also show results in a later chapter obtained with nonlinear non-perturbative models where non-perturbative effects on the mode structure will be evident. These effects are become likelier to be seen the further the modes are far from the linear stability threshold (in the later chapter, we observe such effects for $\gamma/\omega \sim 3\%$). We note that the workflows used in the previous section do not include such features, however they are sufficiently flexible to allow us to study the effect of non-self-consistent mode structure modification. We note that chirping using the HAGIS model has been previously demonstated [99].

To this end, the HAGIS code was extended to support a time dependent mode structure in a simple fashion. On the intermediate time scale in the code, the following procedure is performed. For each mode i, we calculate the energy of mode i; then we update the mode structure of mode i, in this case by updating the parameters according to an analytical expression; then we calculate the energy of mode i after the change in the structure, and we rescale the mode amplitude to keep the energy in the mode constant; finally, if applicable, we proceed to the next mode $i + 1$.

Moving away from the standard ITER 15MA scenario, we perform an exercise based upon the HAGIS/FAC ITER benchmark outlined in reference [65]. We simplify this somewhat by taking a single toroidal mode, and keeping only a single poloidal harmonic, which we parameterize by fitting a radial Gaussian function, using the expression in equation 4.1, with μ and σ representing the mode radial position and the double square

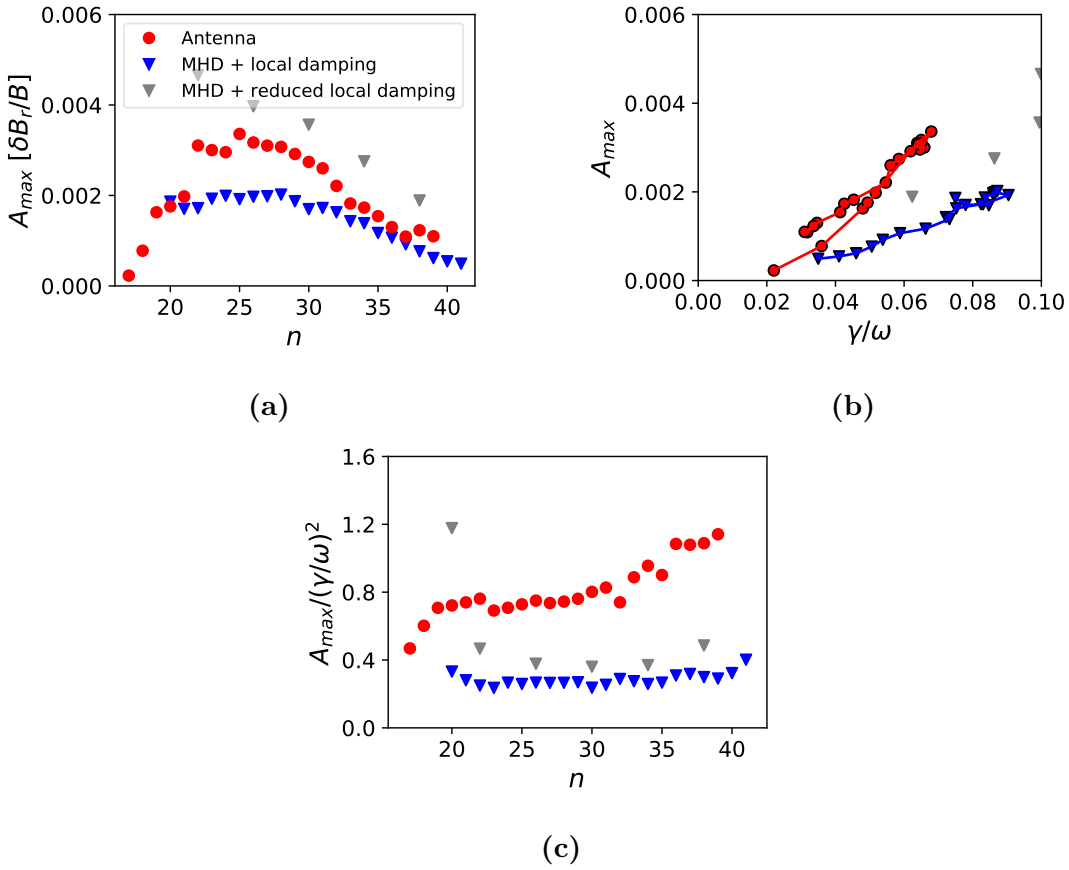

Figure 4.5: Mode saturation amplitude from nonlinear simulations plotted as a function of toroidal mode number (a) and normalized linear growth rate (b) obtained using both the global MHD and global kinetic workflows. Additionally plotted are also points with significantly reduced (by factor ~ 6) damping rates, to demonstrate the sensitivity of the mode saturation to exaggerated uncertainties in the damping rate. In (c), we plot the coefficient calculated by normalizing the saturation amplitude to the square of the normalized growth rate as a function of the toroidal mode number.

width respectively.

$$A(r) = A \cdot \exp\left(-\frac{(r-\mu)^2}{\sigma}\right) \tag{4.1}$$

We then define a time dependent function for the width term, a smoothed step function, as in equation 4.2.

$$\sigma(t) = \sigma_0 + \Delta\sigma \cdot \left(\frac{1}{\pi}\arctan\left(\frac{t-t'}{\Delta t}\right) + 1/2\right) \tag{4.2}$$

We then perform the following experiment, for various values of $\Delta\sigma$, Δt and t', we perform nonlinear simulations and measure the saturation amplitude. We choose parameters according to all combinations of, $\Delta\sigma = \{0.4, 0.2, -0.2, -0,4\}\sigma_0$, $\Delta t = \{1, 5, 10\} \cdot 10^5$, $t' = \{1.0, 1.5, 2.0, 2.5\} \cdot 10^6$, with t in units of the wave period. Some examples of these are plotted in figure 4.6.

The results of this experiment are plotted in figure 4.7. The simplest conclusion is drawn by looking at the solid line. This tells us that the saturation amplitude does indeed weakly depend on the mode width, decreasing by $\approx 25\%$ for $\sigma = \sigma_0/2$, and increasing by $\approx 5\%$ for $\sigma = \sigma_0 \cdot 3/2$. The next conclusion is that most points lie close to the solid line. This implies that the saturation amplitude is most strongly affected by the width at the time during the saturation process. We also note that four cases deviate significantly from the solid line. We also note that these points all correspond to $t' = 2.5 \cdot 10^6$ (the latest time), and all correspond to $\Delta t = 10^5$ (the fastest transition). We also note that the other points with furthest deviation from the line correspond to late times and fast transitions. There are two possible explanations for the deviation. Firstly, from the lower panel of figure 4.6, we see that the deviation from the linear phase has clearly started before $t = 2 \cdot 10^6$, therefore any case where the transition is centered around $t = 2.5 \cdot 10^6$, especially cases where the transition takes place quickly, will have already begun the nonlinear saturation process before any significant modification to the mode structure has taken place. A second possibility is that these fast transitions are too fast for the intermediate time scale implemented in the model. This seems unlikely as the earlier time fast transition cases are closer to the line, however it should be reminded that we apply a piecewise constant width, which only approximates the smooth function in equation 4.2. This introduces some additional uncertainty in the width at the saturation time, therefore the actual width may have been narrower than reported. However, as this is such a large, rapid transition of the mode structure relative to the saturation time scale, we do not consider this to be worth pursuing in more detail.

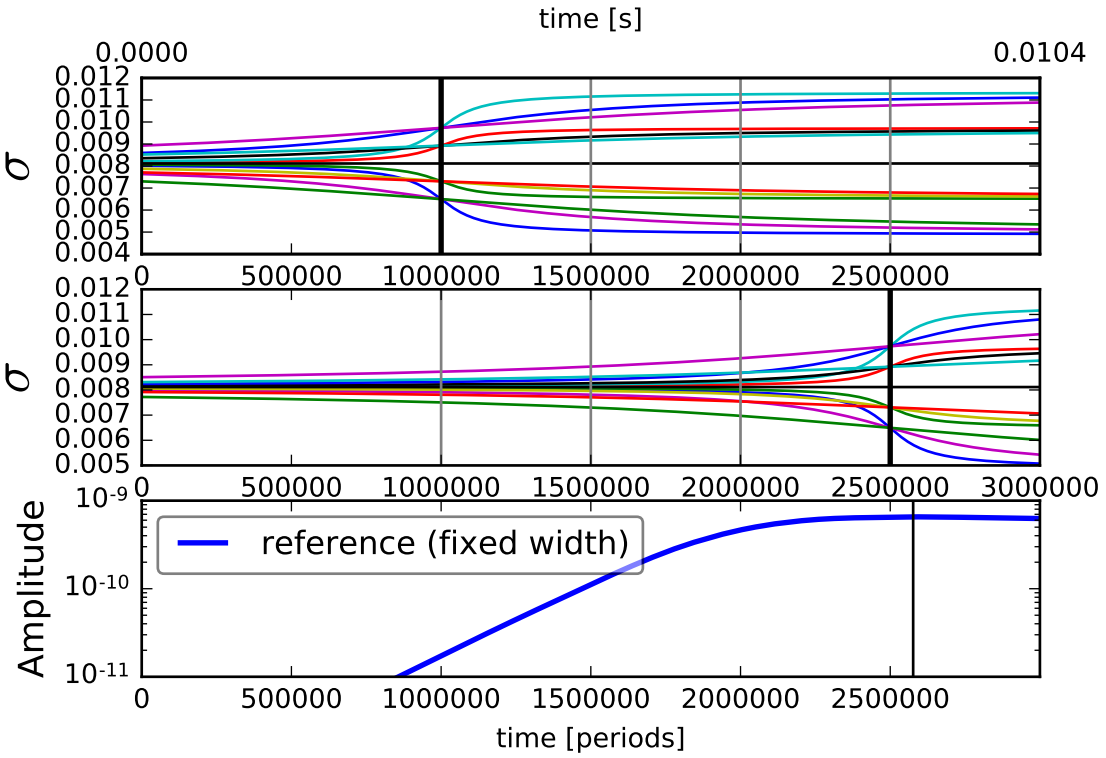

Figure 4.6: In the upper and middle panels, some example functions of $\sigma(t)$ for $t' = 10^6$ (upper) and $t' = 2.5 \cdot 10^6$ (middle). In the lower panel, we plot the amplitude evolution for the simulation with the reference mode width ($\sigma(t) = \sigma_0$), noting that the saturation occurs at $t = 2.58 \cdot 10^6$ (marked with the vertical line).

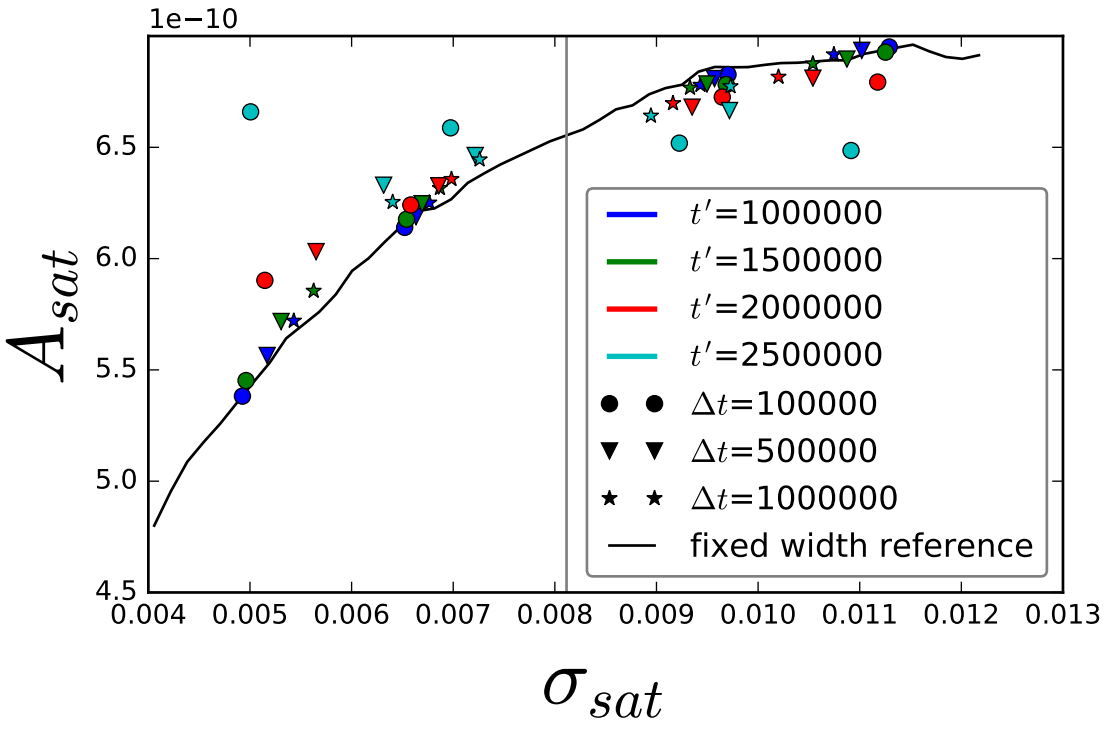

Figure 4.7: Mode saturation amplitude plotted against $\sigma(t_{\mathrm{sat}})$ for 48 combinations of parameters in the time dependent width equation 4.2. The point type denotes the value of Δt and the point colour denotes the value of t'. The solid line shows the values of saturation amplitude obtained when performing a scan of mode widths σ (constant in time). The vertical line corresponds to the nominal case $\sigma(t) = \sigma_0$, the value shown in the lower panel of figure 4.6.

5 Initial value global gyrokinetic simulations of energetic particle driven modes in ITER

5.1 Starting with a reduction of the problem

In order to start simulating the TAE case of ITER, we start by simplifying the problem. We shall then assess the effect of the simplifications and, if necessary, move in the direction of the more complete case. These simplifications have been performed in the interest of numerical stability; numerical cost; to aid comparisons; or to simplify the setup of the study.

The first simplification we perform is to treat the alpha particles as a Maxwell-Boltzmann distributed species with temperature of 900 keV [22, 96]. The next simplification is that instead of treating the background plasma ions as a half-half mixture of deuterium and tritium ions, we can instead treat the background plasma ions as a single species of a hypothetical hybrid isotope, with atomic mass 2.5. The effect of this approach was quantified in [8]. Although the mass ratio between $^{2.5}$DT ions and electrons is approximately 4500, if one increases the mass of electrons in simulations, which affords a numerical saving, the physical results do not necessarily change. Therefore, we start by considering heavy electrons, giving a reduced mass ratio of 200. We also neglect impurity ion species, such as (cold) helium ash and beryllium. The effect of these species is also studied in [8] and will be discussed later. Even though collisional processes are present in plasmas, the physical processes we are investigating are principally collisionless and short in time scale. Therefore, we neglect any collisions, as well as any particle and energy sources and sinks. For long time scales, of the order of the plasma transport time scales, these effects would need to be included. We are interested in core-plasma phenomena, however in this ITER scenario, the background plasma density profiles include large temperature and density gradients in the edge. These exceptionally steep gradients cause problems for the model, which was not intended for studying such profiles. Therefore, we start by considering flattened profiles for the background plasma, taking values for the electron and ion temperatures approximately

equal to their values at the mid-radius. We shall then later keep the effect of the profile variation by removing the gradients only from the steep region of the profiles in the edge. As we are interested in the discrete toroidal mode number eigenmodes, we set the filter in each simulation to include only a single n (unless explicitly stated otherwise), and the corresponding band of poloidal harmonics according to the field-aligned filter $m(r) = \lfloor n \cdot q(r) \rfloor \pm \Delta m$, with Δm chosen as 5. Finally, since a gyrokinetic code is used, it is possible to neglect the averaging of the gyromotion and consider a species as if drift-kinetic. This is always done in the case of the electrons, but can also be done in the case of ions. This slightly reduces the cost of simulations, but also increases the growth rate of the mode and also allows a comparison with hybrid drift-kinetic results.

From hybrid modelling, we expect the TAE drive to be largest in the range of $n = 20$–30. From eigenvalue calculations, we also expect that for $n \gtrsim 20$, the TAE structure should become localized. Also for larger mode numbers, both toroidally and poloidally, we will require more points to accurately represent such a perturbation. When we take all of these factors together, we shall start by studying mode numbers between 20 and 30, starting with the $n = 24$ TAE.

5.2 Initial simulations of n=24 AEs

To initialize a simulation, it is necessary to apply a small initial perturbation. Here, one has several choices as to the nature of such a perturbation. Most of the simulations presented here will be initialized by applying some knowledge of the TAE physics. In this case, we expect that the mode structure primarily comprises two relatively narrow poloidal harmonics (in this case, $m = (n, n+1)$), only narrowly separated, and centred at the location where the safety factor is approximately $q = \frac{m+1/2}{n}$, which in this case is $q = \frac{24+1/2}{24} \approx 1.02$, corresponding to a radial position of $s \approx 0.4$. We therefore initialize two radial Gaussian perturbations with the corresponding m and n. This gives the same final results as an alternative initialization, but enables the simulation to establish the linear mode structure faster. An alternative initialization is used when we do not have a good estimation of the linear mode structure, which has a much broader radial envelope, and initializes many poloidal harmonics according to the local safety factor. Finally, we can also initialize an initial perturbation with a random radial structure, according to a range of poloidal mode numbers. This is not conducive to obtaining clean, well defined modes, but can be useful for probing the continuum spectrum.

In figure 5.1, we show snapshots of poloidal cross sections of the simulation domain, showing the electrostatic and electromagnetic potential perturbations due to the mode obtained via a simulation of the $n = 24$ TAE mode structure with double the nominal EP density. This simulation has been done according to the most simplified case

above. We see that the two quantities show broadly similar shapes, but whereas the electrostatic potential shows a maximum on the low field side (right of the figure), the electromagnetic potential has a maximum on the high field side (left). This observation is in line with analytical expectations. To facilitate more detailed analysis and comparisons of mode structures however, it is convenient to look at the radial structures of the different poloidal harmonics of the mode. We therefore show in figure 5.2 the poloidal decomposition of figure 5.1a, plotting the absolute value of each harmonic. We typically only plot harmonics whose peak is at least 5% of the peak harmonic. This forms our primary method for evaluating the spatial structure of the modes. By also looking at the time evolution of the electrostatic potential (taking points on the low field side at poloidal angle of 0), we can measure the frequency and growth rate of the mode. This analysis is shown in figure 5.3.

We here take the opportunity to introduce the diagnostic method presented here for the frequency and growth rate analysis. In the upper part of figure 5.3, we see the signal envelope. This signal envelope is in fact the magnitude of the n^{th} component of the toroidal Fourier component of the signal. The growth rate of this envelope is fitted in the shaded region, giving the growth rate of the mode. The signal (at a fixed point in space) is then divided by the exponential fit of the envelope. This normalized signal is then Fourier transformed in time, shown in the lower part of the figure. The peak of the Fourier transform is then fitted to give the real frequency of the mode. In general, this procedure is performed at all radial positions independently. Figure 5.4 shows the measurements of the growth rate and real frequency at all radial positions. The fitting of the envelope can be performed using different methods. Here, we take the fit obtained by taking the points which are local extrema. The fit in figure 5.3 is evaluated at the position of the maximum of the mode and corresponds to $s = 0.4$ in figure 5.4b.

If we simulate instead the full radial domain, we can test the validity of running this simulation on a reduced annulus. The mode structure, shown in figure 5.5, matches closely that of the annular simulation, as do the frequency and growth rate, which are shown in figure 5.6.

We can also run the same simulation without energetic particles present. By doing this, we can see the initial perturbation decay, showing us two interesting phenomena. To do this, we use the alternative initialization described at the beginning of this section. Firstly, we note that we no longer only see the most strongly driven mode, but now a mixture of all the weakly damped modes with $n = 24$. By looking at the radial frequency plot in figure 5.7a, we see that also other modes such as the ellipticity induced Alfvén eigenmodes (EAEs) and triangularity induced Alfvén eigenmodes (NAEs) are present at higher frequency. The second feature is that as the initial perturbation is decaying away, we also see in the radial frequency spectrum the nature of the continuum.

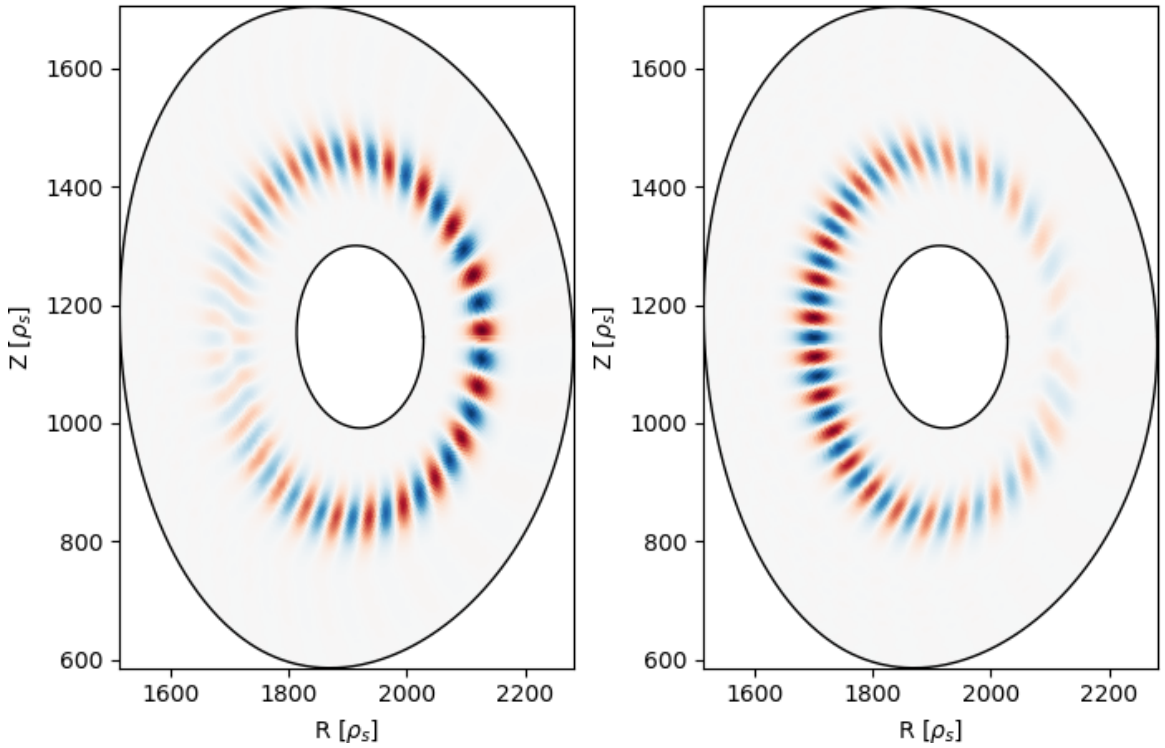

(a) Perturbed electrostatic potential. **(b)** Perturbed electromagnetic potential.

Figure 5.1: (a) Electrostatic and (b) electromagnetic potential perturbations in the poloidal plane, showing $n = 24$ TAE simulated on an annulus with reduced mass ratio.

We include in this figure also the Alfvén continuum calculated using both the ideal MHD and the kinetic versions of the eigenvalue solver LIGKA. Figure 5.7b shows the same exercise, but with the noisy, random, initialization also described at the beginning of this section. In this case, we can see the nature of the continuum in both cases, although the modes are more clearly visible in figure 5.7a. In that case, we plot the radial harmonic structure in figure 5.8. Here, we can see that the radial harmonic structure is a superposition of different modes, demonstrating that weakly damped sub-dominant stable modes exist.

5.2.1 Ellipticity induced Alfvén eigenmode (EAE)

In simulations with no energetic particles, for example as shown in figure 5.7, we can observe the $n = 24$, $m = (24, 26)$ EAE mode at a higher frequency ($\omega \approx \omega_A$) and at a slightly larger radial position ($s \approx 0.52$) as compared to the TAE. In figure 5.8, we can see the corresponding peak of the mode structure. In order to study this mode further,

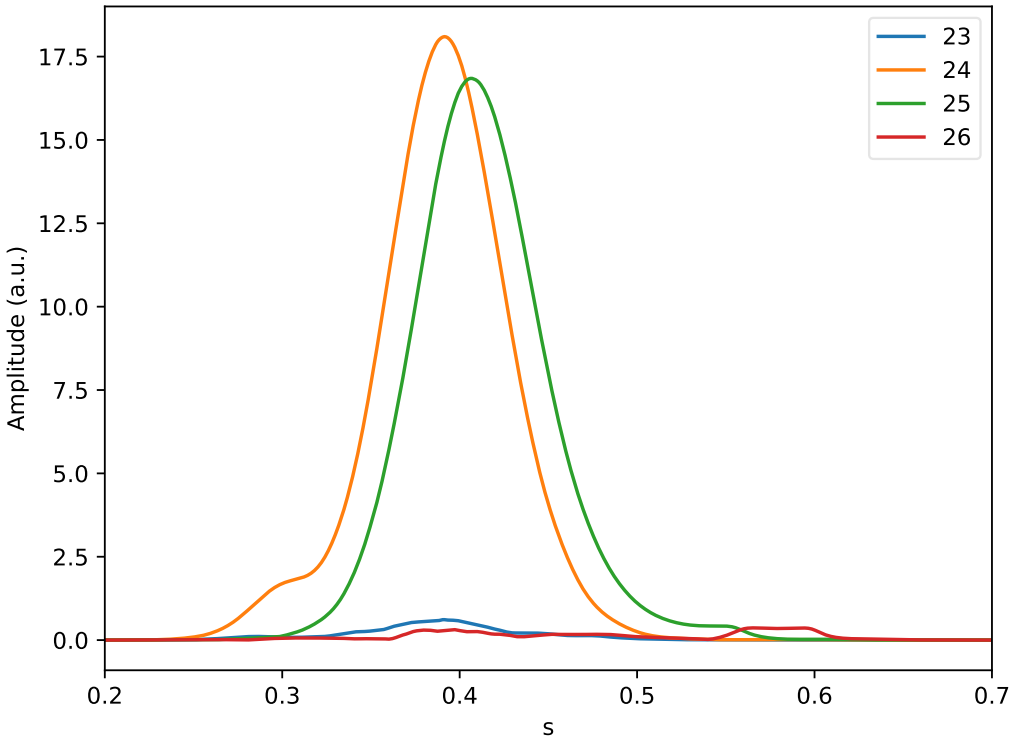

Figure 5.2: Mode structure for an $n = 24$ TAE simulated with flat background profiles, double the nominal EP density, reduced mass ratio of $m_\mathrm{i}/m_\mathrm{e} = 200$, simulated on an annulus of $s = 0.2$–0.7. The labels refer to different poloidal harmonics of the electrostatic potential.

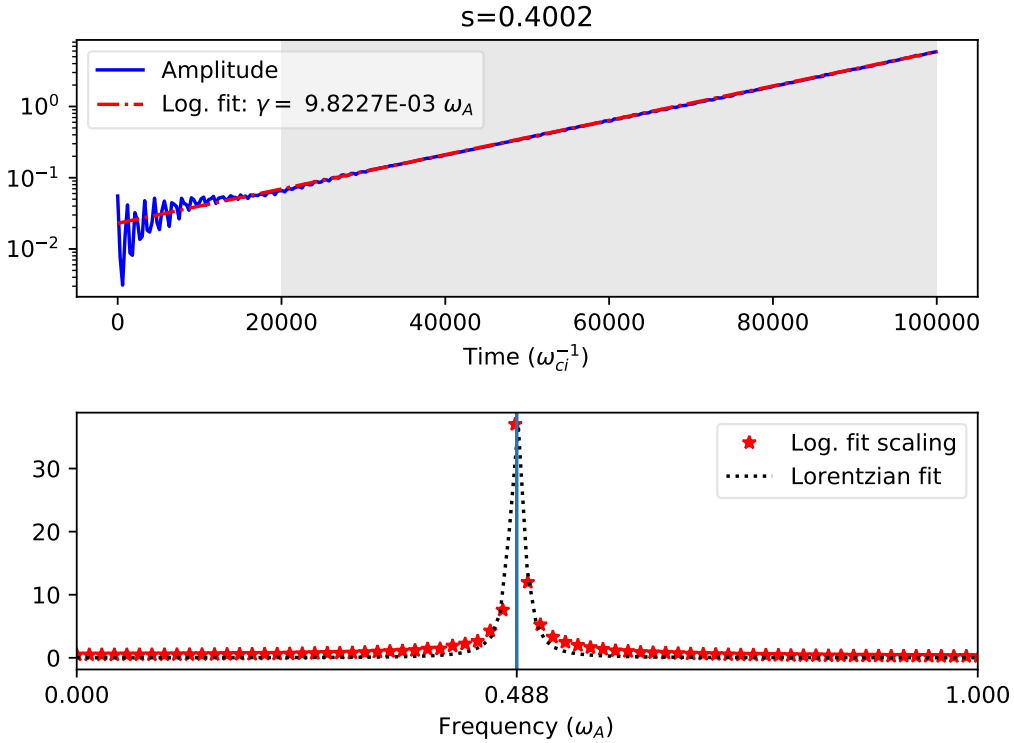

Figure 5.3: Frequency and growth rate analysis performed for the same $n = 24$ simulation as shown in figure 5.2. In the upper part of the figure, we show the envelope evolution and the fit for the growth rate, which is written in the label of the figure. In the lower part, we plot the amplitude of each frequency bin (in units of the Alfvén frequency) of the normalized signal after applying a Fourier transform. For the purpose of interpolating the frequency, we also fit a Lorentzian function to the values in the bins of the Fourier transform.

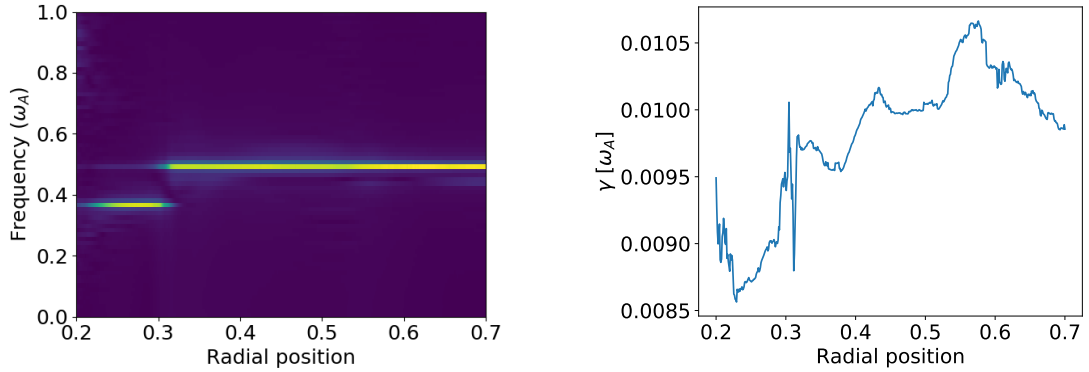

(a) **(b)**

Figure 5.4: Measurements of the frequency and growth rate at different radial positions, applying the same procedures as in figure 5.3, which equates to a cut of these figures at radial position $s = 0.4$. On the left, we plot the frequency spectrum, noting that each radial position is normalized independently based on its fit in the right figure.

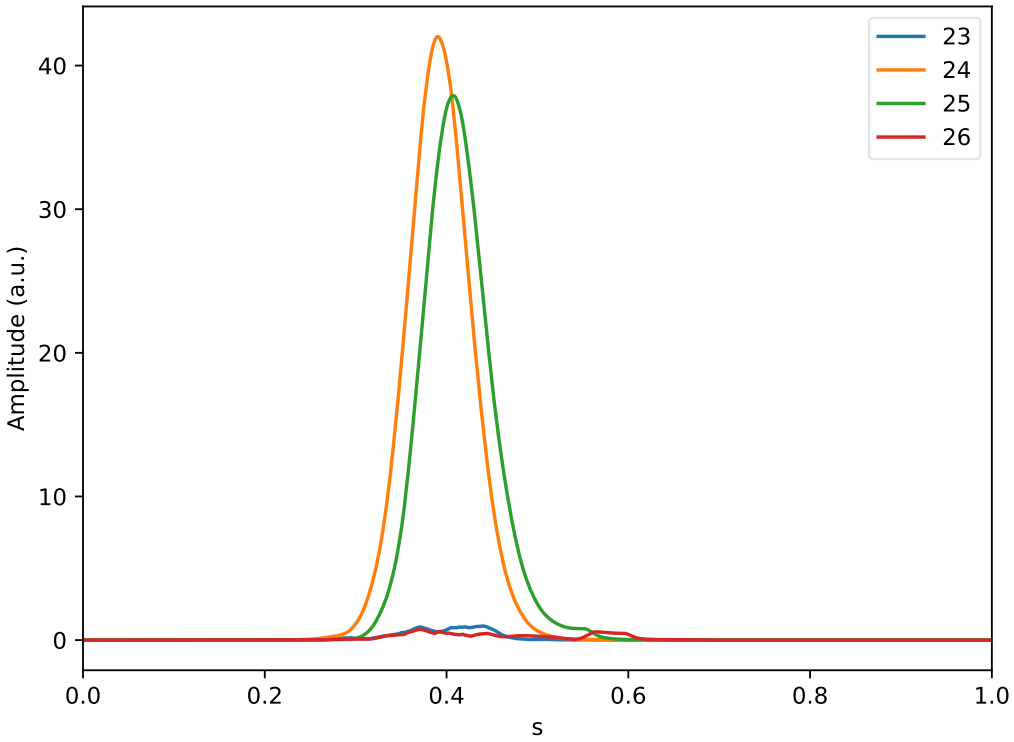

Figure 5.5: Mode structure for an $n = 24$ TAE simulated with flat background profiles, double the nominal EP density, reduced mass ratio of $m_i/m_e = 200$, simulated on the entire radial domain. The labels refer to different poloidal harmonics of the electrostatic potential.

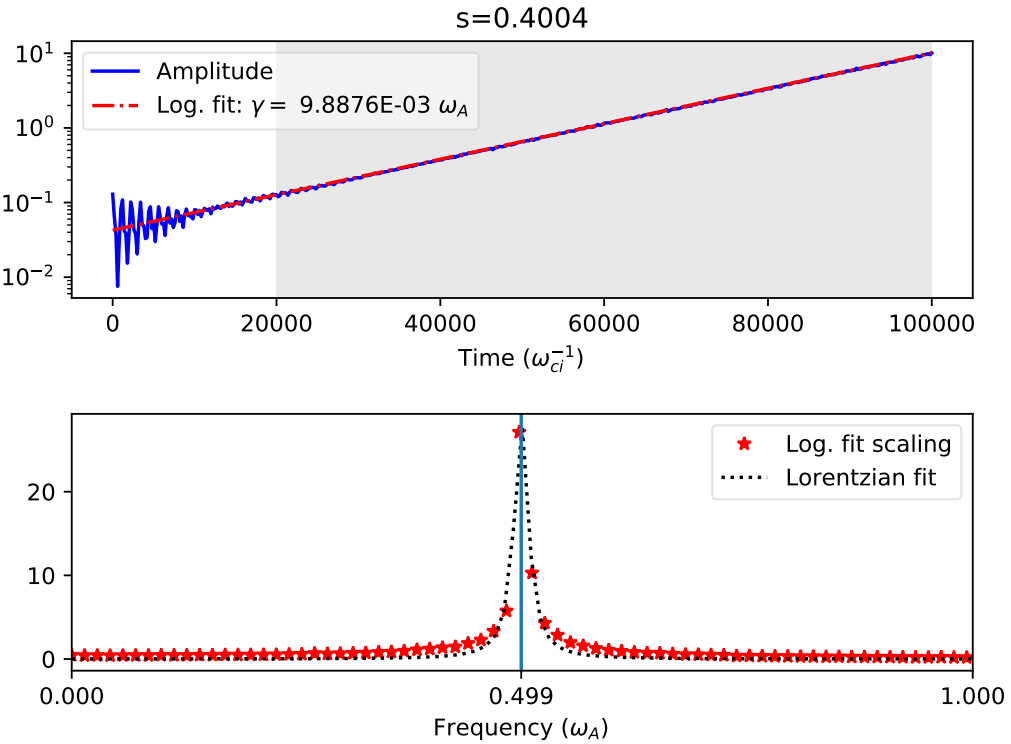

Figure 5.6: Frequency and growth rate analysis performed for the same $n = 24$ simulation as shown in figure 5.5.

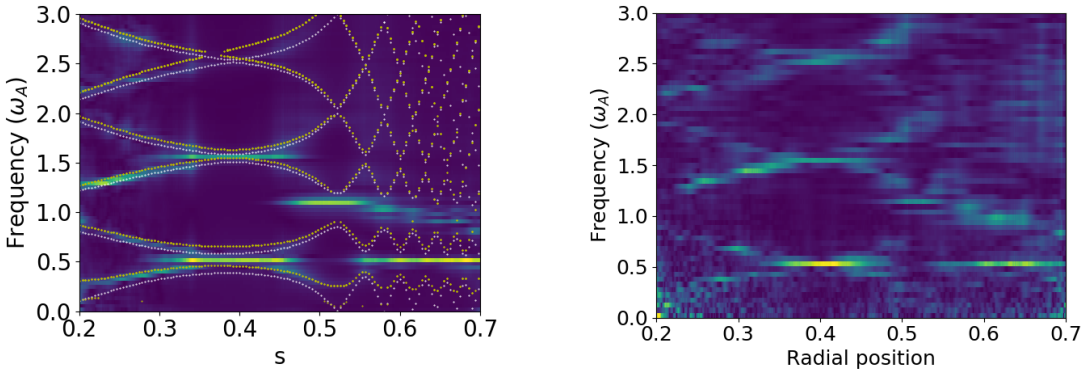

(a) **(b)**

Figure 5.7: Radial frequency spectrum for $n = 24$ simulated with zero EP density. Initialization is based on (a) broad radial structure and poloidal mode numbers chosen to be field aligned and (b) random radial structure and a fixed range of poloidal mode numbers. On the left, we overplot the Alfvén continuum from the ideal MHD version of LIGKA in white and the Alfvén continuum from the analytical kinetic version of LIGKA in yellow.

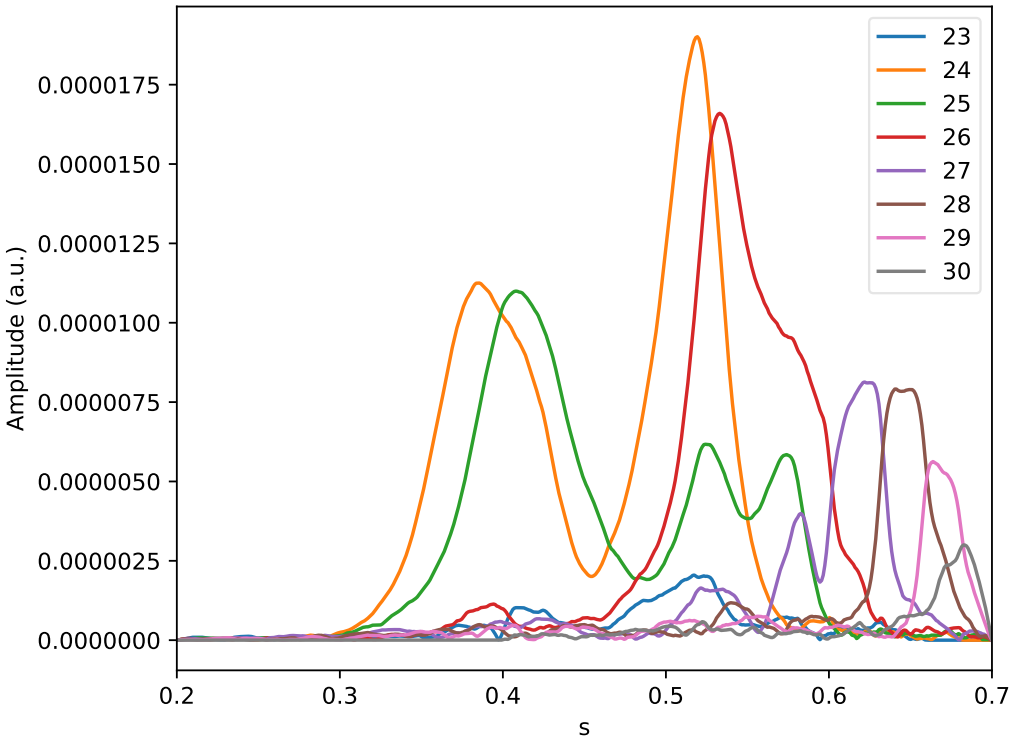

Figure 5.8: Radial harmonic structure for $n = 24$ simulation without EPs present. The simulation corresponds to figure 5.7a.

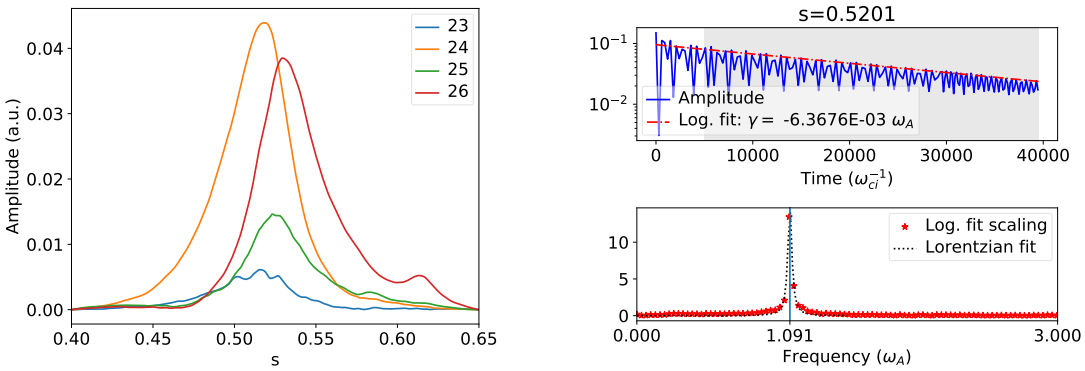

(a) **(b)**

Figure 5.9: Radial harmonic structure (left) and damping and frequency measurements (right) for an $n = 24$ EAE simulation without EPs present.

it is possible to create a simulation specifically tailored to this mode. We initialize now a narrower envelope with harmonics $m = 24$ and 26 located close to the $q = 25/24$ rational surface, as expected from analytical theory, and we can further decrease the simulation domain for this exercise given the locality of the EAE eigenfunction to help us to exclude other modes. In figure 5.9a, we show the radial mode structure of this EAE, with the frequency and damping measured in figure 5.9b. With a damping rate of only 0.6% of the mode's real frequency, this mode is only weakly damped. However, as we only see this mode in simulations without energetic particles present, it is reasonable to infer that these modes are at most weakly driven by the energetic particles we consider. This conclusion is in agreement with previous studies of EAEs [61].

5.2.2 Odd-parity TAEs

Within the TAE gap, there can exist eigenmodes with frequencies close to the lower part of the gap, in which case the poloidal harmonics of the electrostatic potential have the same sign and are said to be even-parity, or having a ballooning structure. This includes the TAE modes found so far in this section. In addition to these modes, there can also be eigenmodes with frequencies close to the top of the gap, which lead to odd-parity modes, which possess anti-ballooning structures. Here, we modify slightly the initialization, by multiplying the upper harmonic of the TAE by a prefactor of -1. If we then perform this simulation without energetic particles, then after some time, we can see both the even and the odd parity TAEs. In figure 5.10a, we see the splitting of the frequency close to $\omega_A/2$ into the even TAE and the odd TAE at the lower and higher frequencies respectively. An analysis of the frequency close to the mode position

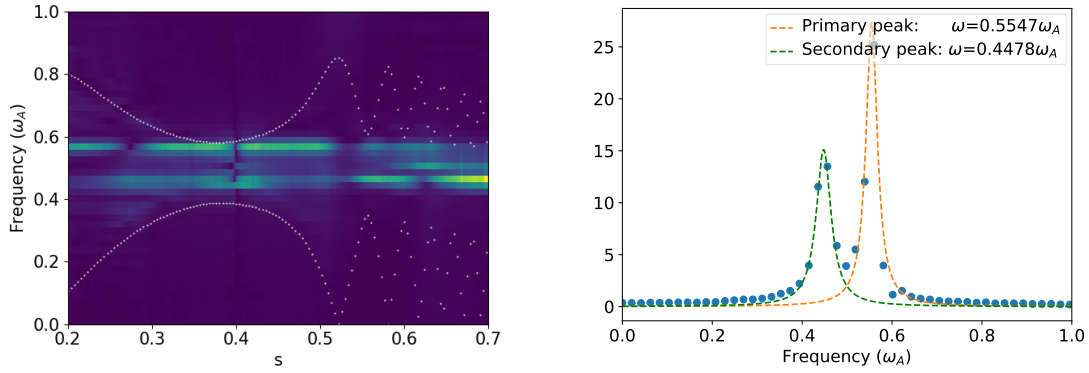

(a) **(b)**

Figure 5.10: Analysis of the frequency for $n = 24$ TAEs without energetic particles, initialized with odd parity. In (a), the radial frequency spectrum, with the Alfvén continuum, as calculated from the ideal MHD version of LIGKA, overplotted in white. In (b), the frequency analysis performed at $s = 0.35$ showing the fit of the twin frequency peaks.

(at $s = 0.35$) reveals peaks of the frequency at 0.55 and 0.45 ω_{A} (which is shown in figure 5.10b). In the figure, we overplot the ideal Alfvén continuum obtained from the ideal MHD version of LIGKA, and we see that, as expected, the mode frequencies are close to the upper and lower extrema of the TAE gap. We note by their absence in simulations with EPs present, although we see that they are not heavily damped, that the odd-parity TAEs are not effectively driven by the EPs considered.

5.3 Dependence on the electron mass

As previously mentioned, it is convenient to increase the electron mass, (in other words reduce the mass ratio (m.r. $= m_{\mathrm{i}}/m_{\mathrm{e}}$)), in order to reduce the difficult and cost of initial value simulations. For a fixed energy, the electron velocities scale inversely with the square root of the electron mass. The increased cost associated with large mass ratios comes due to the reduction on the time step (dt) needed to resolve the faster electron motion. If there were to be no further change in the plasma dynamics, we would therefore expect to require a time step which scales with the square root of the electron mass.

In addition to the numerical considerations outlined above, there is also a physical consequence of changing the electron mass. As one of the physical damping mechanisms of the mode, electron Landau damping, is dependent on velocity space resonances, typically located close to the trapped-passing boundary of the thermal electrons. Therefore, we expect this term to be affected by any changes to the electon mass. From

previous analytical theory (e.g. equation 34 of reference [74]), we can estimate that the electron Landau damping should be proportional to v_A/v_e, and therefore proportional to $m_e^{1/2}$. We expect, therefore, that for an artificially large electron mass, we would see an increased electron Landau damping. However, in the presence also of other damping contributions, some reduction in the mass ratio should not significantly affect the mode's damping rate.

To estimate the damping of a mode using an initial value method, we can remove the drive (i.e. setting the EP density to zero), and measure the growth rate. However, there are a number of issues with such a method which complicate measuring the growth rate of decaying simulations. Firstly, there is no guarantee that the observed mode (assuming that we observe a single mode) is exactly the same mode as in the driven case. Incidentally, this statement can also be applied to cases with a significant increase in the EP density. Secondly, there is an additional numerical problem associated with accurately measuring strongly decaying modes. As the amplitude of the mode decreases, the signal to noise ratio decreases, decreasing the quality of the simulation. Therefore, particularly if we want to exclude the initial phase, fitting only the time frame with a well established eigenfunction, then we are typically left to only a relatively short time window for analysis. An alternative method, which can help in such cases is to measure the damping as the offset to the linear relation between EP density (i.e. drive) and growth rate. This method of course can only be used when the energetic particles' contribution to the eigenmode is non-perturbative.

Here, we run a series of $n = 26$ TAEs, much like the $n = 24$ examples earlier in this section. In this case, to aid quantitive comparisons, all simulations are run with the equilibrium scaled such that the on-axis value of the safety factor, q_0, is 0.985 (previous results in this section used the very similar value of 0.99). All of the following cases use the experimental profiles, simulated mostly in annulus corresponding to a radial domain of $s = 0.2$ to 0.7. We perform these simulations without EPs, with the nominal EP density, and with double the nominal EP density, performing simulations with ion-to-electron mass ratios ranging from very low (50), up past realistic (4557), up to a few examples with a super-realistic (8000) mass ratio. For each case, we also perform a convergence test with respect to the time step, in order to see the effect of the electron mass on the time step convergence.

For each of these cases, we perform the following analyses. We fit the growth rate of each case, by fitting to the absolute values of extrema point of the 26th toroidal fourier component of the potential, measured at $s = 0.453$ (the location of $q = \frac{26+1/2}{26}$). We also perform a frequency analysis, by taking the temporal fourier transform of the potential at a single point ($\phi = \chi = 0, s = 0.453$), normalized to the previously fit growth rate. This frequency and growth fitting is performed within a time window, namely $t = 0$ to 20000 ω_{ci} for the cases without EPs, and $t = 20000$ to 40000 ω_{ci} for

the cases with EPs. Within such a window, it can be difficult, particularly in decaying cases to measure the growth (decay) rate, but as we perform the same analyses for all mass ratio and time step combinations, we can still perform a fair relative comparison. Additionally, we shall perform a more detailed analysis for an individual example, which allows us to put the decaying results in the context of an absolute comparison. As a secondary method for measuring the growth rate, we perform the fit not on the extrema, but on all values. By comparing the growth rate measured with the primary and secondary methods, we can approximate an uncertainty range of the fit of our measured growth rate, which is shown in figure 5.12. We observe that the decaying cases have a non-negligible error, understandable because of the shorter fit range, the weaker signal, and the coexistence of multiple decaying oscillations.

The results of these studies are shown in figure 5.11. Here, we plot the measured quantities (frequency or growth rate) against the time step, normalized by multiplying by the mass ratio.

From these results, we draw the following conclusions. Firstly, as we increase the ion to electron mass ratio, we must decrease the time step almost proportionally in order to see the same rate of convergence on the measured frequency or growth rate. This is an interesting numerical result as, as we discussed earlier, one may expect that the time step convergence would go only with the square root of the mass ratio. However this expectation was made only if there is no modification to the mechanisms involved. We also know that the spatial scale over which the interaction is taking place is reduced in addition to an increase in the velocity and hypothesize that these effects combine to cause the observed behaviour of the time step. The effect of a change in the size of the spatial scales can be observed by looking at results obtained from an eigenvalue code. We propose that the generality of this observation should be further investigated. Secondly, the physical observation is that the growth rate is reduced in cases with artificially heavy electrons. This is consistent with our previous expectation regarding the increased electron Landau damping. From the figure 5.11a, we note that the lines are approximately parallel, therefore it is a reasonable estimate that the error in the growth rate measured with a well converged mass ratio of 200 would differ from a well converged mass ratio of 4557 by approximately the width between the two lines in the middle of the figure. This would suggest that an extrapolated value for the converged growth rate with mass ratio of 4557 of $\approx 0.02\omega_{\mathrm{A}}$, or $\approx 4\%$ of ω. Combining these two observations, the motivation for using reduced mass ratios is confirmed, however when wanting to do quantitive comparisons, we must choose a sufficiently small electron mass such that the error due to any increased electron Landau damping is small compared to the growth rates. Later results with mass ratio of 200 will use the value for the time step of 1.875, corresponding to the third point from the left of the line labelled “200” of figure 5.11a.

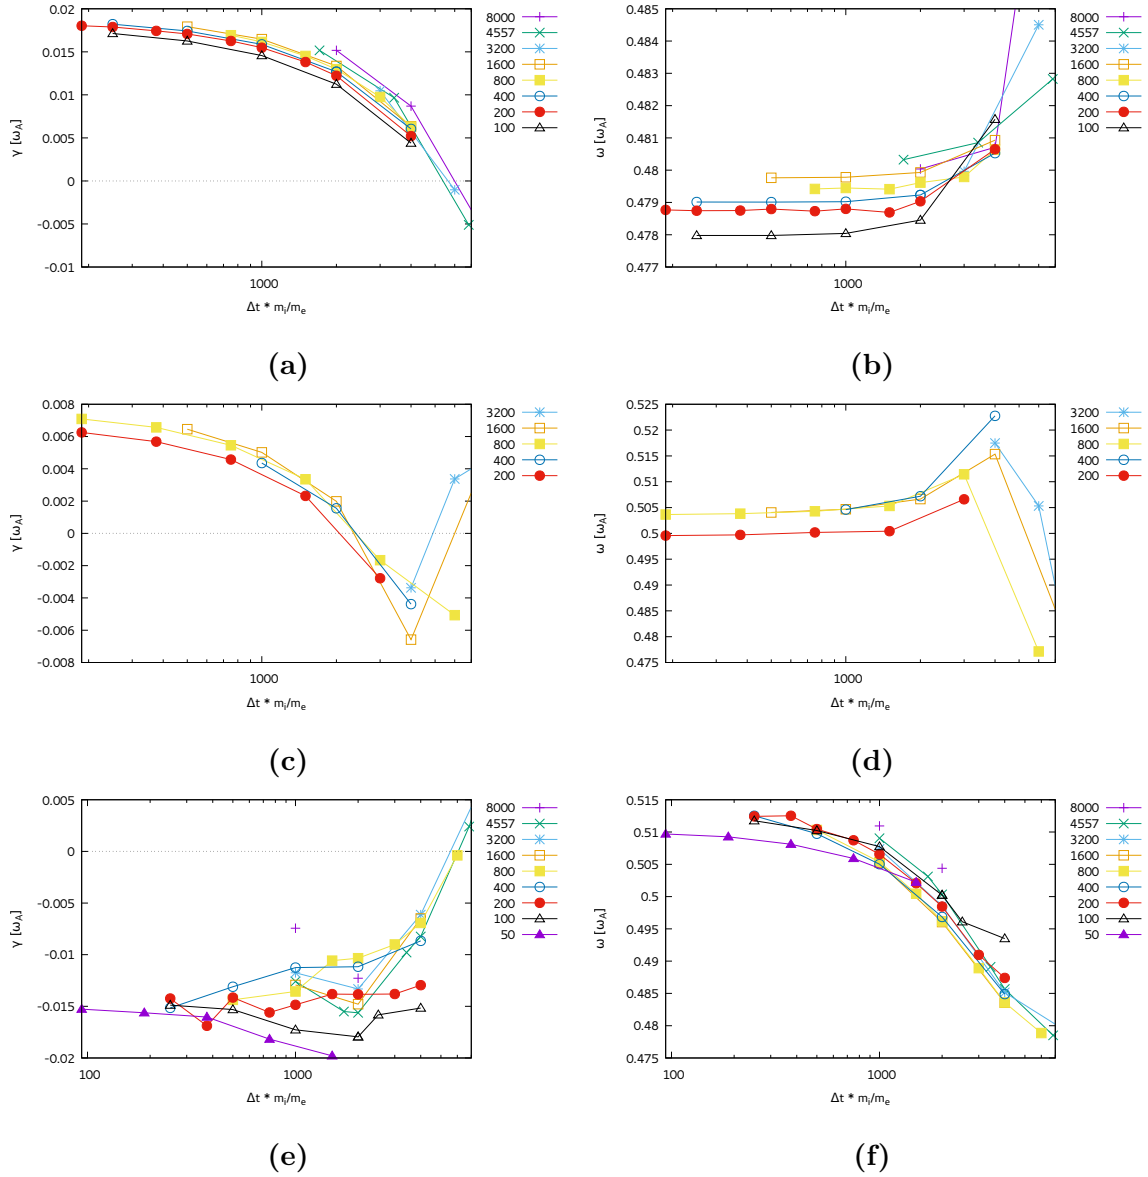

Figure 5.11: Growth rates (a, c, e) and frequencies (b, d, f) taken from simulations of an $n = 26$ TAE run with shaped profiles on an annular ($0.2 \leq s \leq 0.7$) domain for different values of the electron mass and numerical time step. The alpha particle densities are double (a, b), nominal (c, d), or zero (e, f). The legend refers to the ratio between ion and electron mass, which also normalizes the values of the horizontal axis.

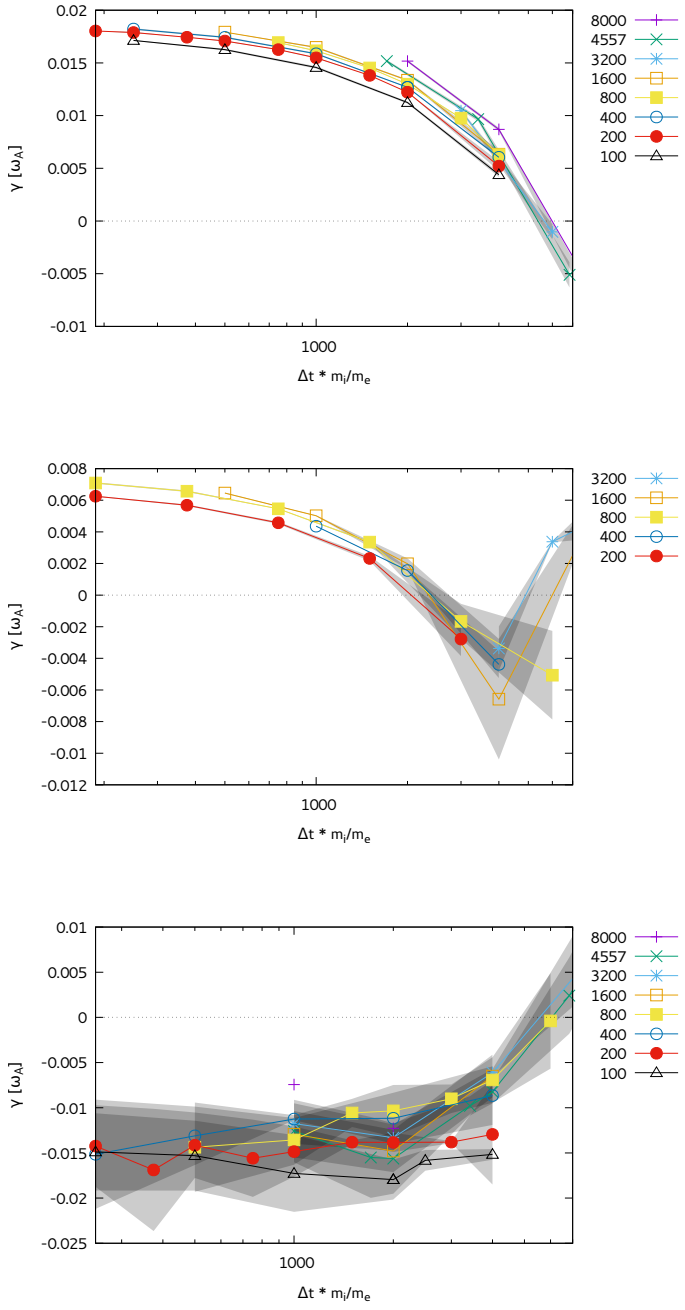

Figure 5.12: The uncertainty on the fit of the growth rates, as per figure 5.11, approximating the uncertainty by taking the discrepancy between measuring the growth rate by fitting the entire signal or by fitting only the extrema of the signal. For double (top), nominal (middle), and zero (bottom) alpha particle density.

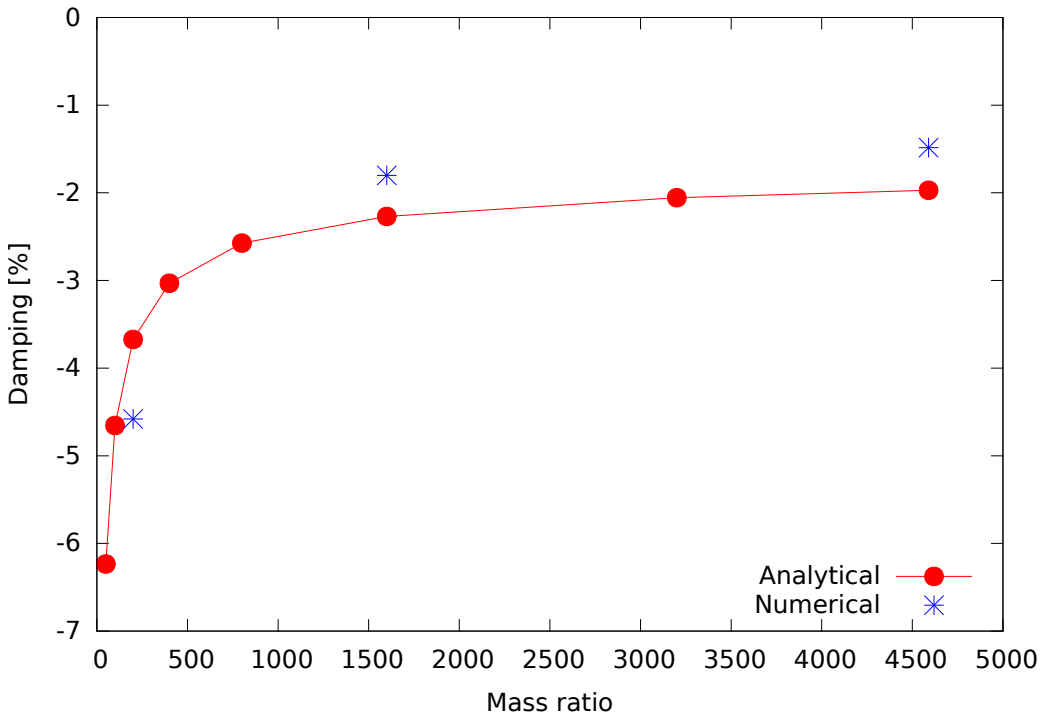

Figure 5.13: Damping rates as a function of m_i/m_e, calculated with LIGKA, measured in units of the mode real frequency. Data range from realistic mass ratios (4557) down to heavy electrons (50). The red points, labelled 'analytical' were calculated using the analytical version of LIGKA (as the rest of this work), whereas the blue points used the full numerical orbits for integrations.

Regarding the cost of the study, to run a simulation to $t = 40000\omega_{\mathrm{ci}}$ requires ≈ 21000 steps with $\Delta t = 1.875$. Each time step has a cost, depending on particle resolution and parallelization, of approximately 0.5 core-hours. However, a similarly converged point with realistic mass ratio would require 486000 steps, or 0.25 million core-hours, which is beyond what can reasonably be performed for a single short linear single-n simulation.

As we compare the frequencies measured in cases without EPs, and with single and double EP density, we observe that the energetic particle pressure gradient reduces the real frequency, pushing the TAE frequency closer to the lower continuum. This result is in line with the expected behaviour [100–102].

For a comparison regarding the dependence of the damping rate on the electron mass, LIGKA was modified to allow the use of unrealistic mass ratios. These results are shown in figure 5.13. Here, we observe qualitatively the same behaviour, although for small mass ratios there are some quantitive discrepancies that are not fully understood. The agreement over a large range close to realistic mass ratios is, however, satisfactory.

We also note that the decay rate of the cases without energetic particles is not the same as the decrease in growth rate between the double and single EP density cases. However, due to the difference in the fitting window, we do not yet infer a conclusion from this. As discussed before, there are some difficulties associated with measuring damping rates from decaying simulations, in particular due to the presence of additional oscillations. Therefore it is difficult to accurately extract the damping rate in the late period of simulations.

5.4 Stability of Toroidal Alfvén Eigenmodes in ITER

Having demonstrated the capability of simulating Alfvén eigenmodes in parameters similar to the those in ITER, we focus on the TAEs in cases with shaped background plasma profiles, cutting only the gradients from the edge, the on-axis safety factor value $q_0 = 0.985$, and looking at the effects of mode globality, the linear spectrum of modes present, and whether nonperturbative features in the modes are apparent. These are evaluated at double EP density.

To this end, we shall focus in particular detail on a few example modes (high n, low n), exhibiting varying behaviours, before summarizing the linear mode spectrum with a scan over toroidal mode number.

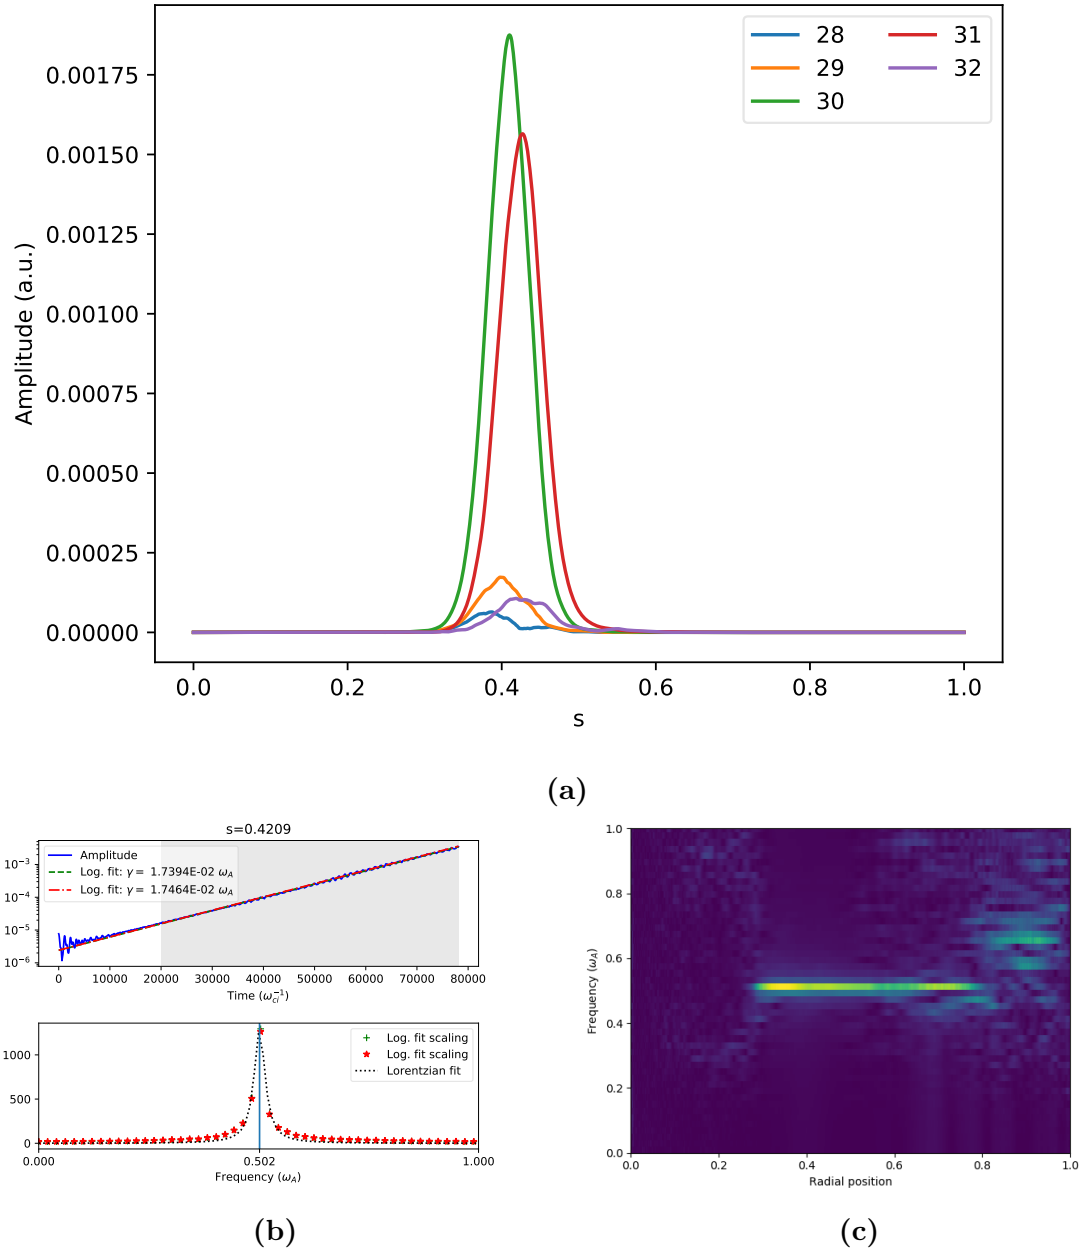

Figure 5.14: $n = 30$ TAE simulated in a global domain. (a) radial structure of poloidal harmonics, (b) growth rate and frequency fit at mode position, (c) radial frequency spectrum of mode.

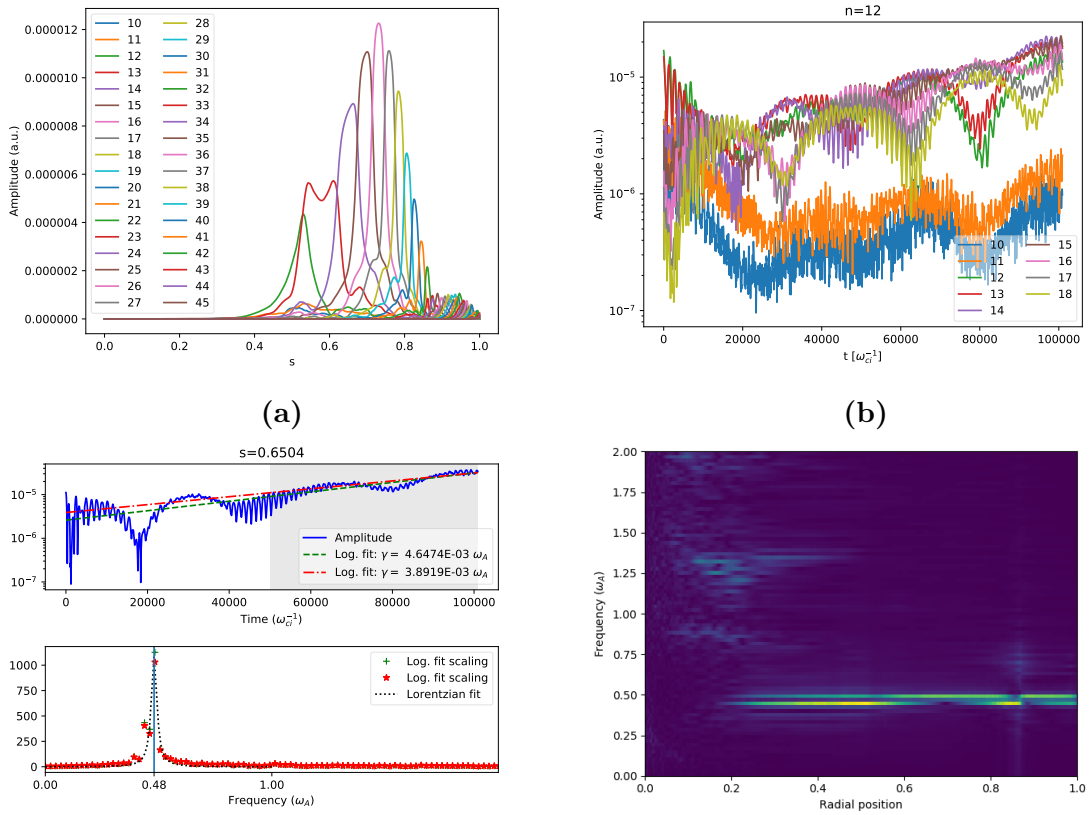

Figure 5.15: $n = 12$ TAEs simulated over a global domain. (a) radial structure of poloidal harmonics, (b) temporal evolution of peak values for each poloidal harmonic, (c) growth rate and frequency fit at mode position, (d) radial frequency spectrum of mode.

5.4.1 High toroidal mode numbers

Our first example is that of a high toroidal mode number, $n = 30$, similar to the $n = 24$ and $n = 26$ cases shown previously. The basic analysis of this mode is shown in figure 5.14. From the plot of the poloidal harmonics, we see that the mode structure is well localized close to $s = 0.42$, the position where $q = \frac{30+1/2}{30}$. Looking at the plot of the mode evolution, we measure a growth rate, $\gamma \sim 0.174\ \omega_A$, and a frequency of $\omega = 0.5$, corresponding to a normalized growth rate, $\gamma/\omega \sim 3.5\%$. From the plot of the frequency spectrum, we see that not only is the mode well localized, but also that it is single frequencied.

5.4.2 Low toroidal mode numbers

Our next example is a case with a lower toroidal mode number, $n = 12$. The analysis of this mode is shown in figure 5.15. In contrast to the case with $n = 30$, this mode has a radially extended mode structure. In figure 5.15b, we observe that the relative peak amplitudes of the harmonics oscillate on a slow time scale as well as growing. In figure 5.15d, we see a splitting of the frequencies present across a large radial domain. If one were to look at the time evolution of the radial structure, one would see a beating pattern of the $m = 12, 13$ peaks close to $s \sim 0.5$ with respect to the outer harmonics. Coupled together, these signatures provide evidence of multiple $n = 12$ TAE modes with global structures and similar frequencies and similar growth rates. Because of the interaction of these modes, it is difficult to obtain an accurate measurement for the growth rates, which are approximately $\gamma/\omega \sim 0.75\% - 1.2\%$. For this case, LIGKA also found three $n = 12$ global TAE eigenmodes with similar frequencies. These eigenfunctions are shown as a comparison in figure 5.16, noting here that the weakly damped outer TAEs are found in both gyrokinetic models. These were found to be important for multi-mode nonlinear treatments in reference [20], and play a significant role in chapter 6.

Noting that this mode exhibits global features, we consider a simulation of $n = 12$ in a reduced annular domain. The results of this are shown in figure 5.17. We note that, in comparison to figure 5.15, the mode structure is cut at $s = 0.7$ according to the boundary condition. Interestingly, the mode growth rate is broadly similar, however there is now no evidence of a beating of different modes, and only a single peak in the frequency is observed.

These observed effects highlight the importance of using global models as well as motivate the retention of linearly subdominant unstable and weakly damped modes in hybrid simulations.

5.4.3 TAE linear spectrum

Now that we have shown some examples of local and global TAE mode structures at low and high toroidal mode numbers, we present results from a scan across a broad range of mode numbers. These results are shown in figure 5.18. We make the following observations. Firstly, we observe a peak in the mode growth rates in the range of toroidal mode number $n \approx 25 - 32$. Secondly, for large mode numbers, where, as discussed previously, the mode structure is relatively well localized, a reduction of the simulation domain has no significant effect on the linear properties of the TAEs. However, for lower mode numbers, we observe larger discrepancies in the linear growth rates, particularly in the range $n = 12 - 18$. Finally, global results for $n = 6$ and $n = 8$

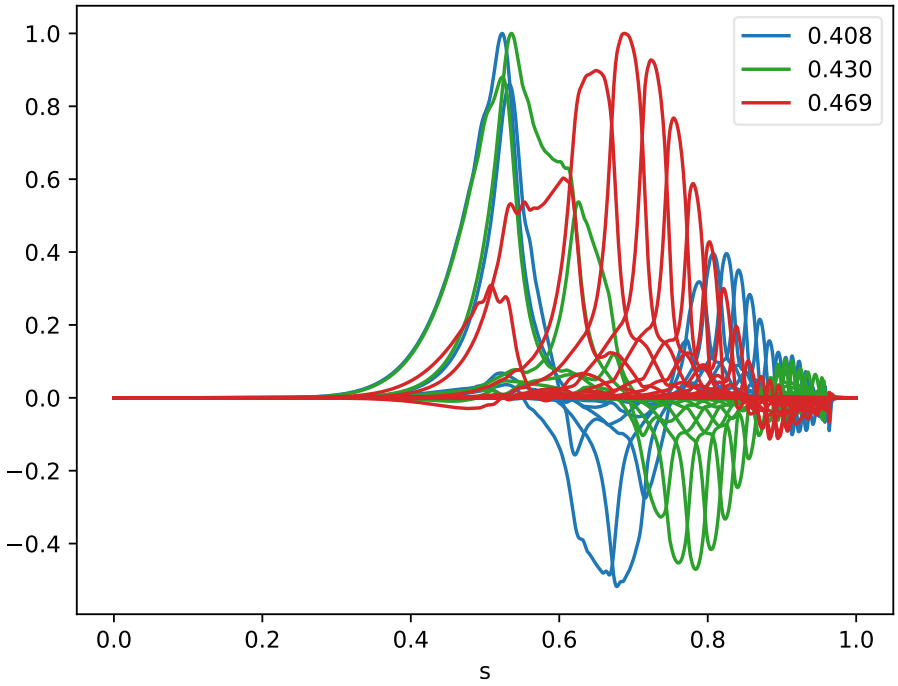

Figure 5.16: Radial harmonic structure for three TAE eigenfunctions found for $n = 12$ with LIGKA, plotting the real part of the electrostatic potential. These eigenfunctions were calculated for use in reference [20]. The legend refers to the frequencies of the modes, normalized to ω_A.

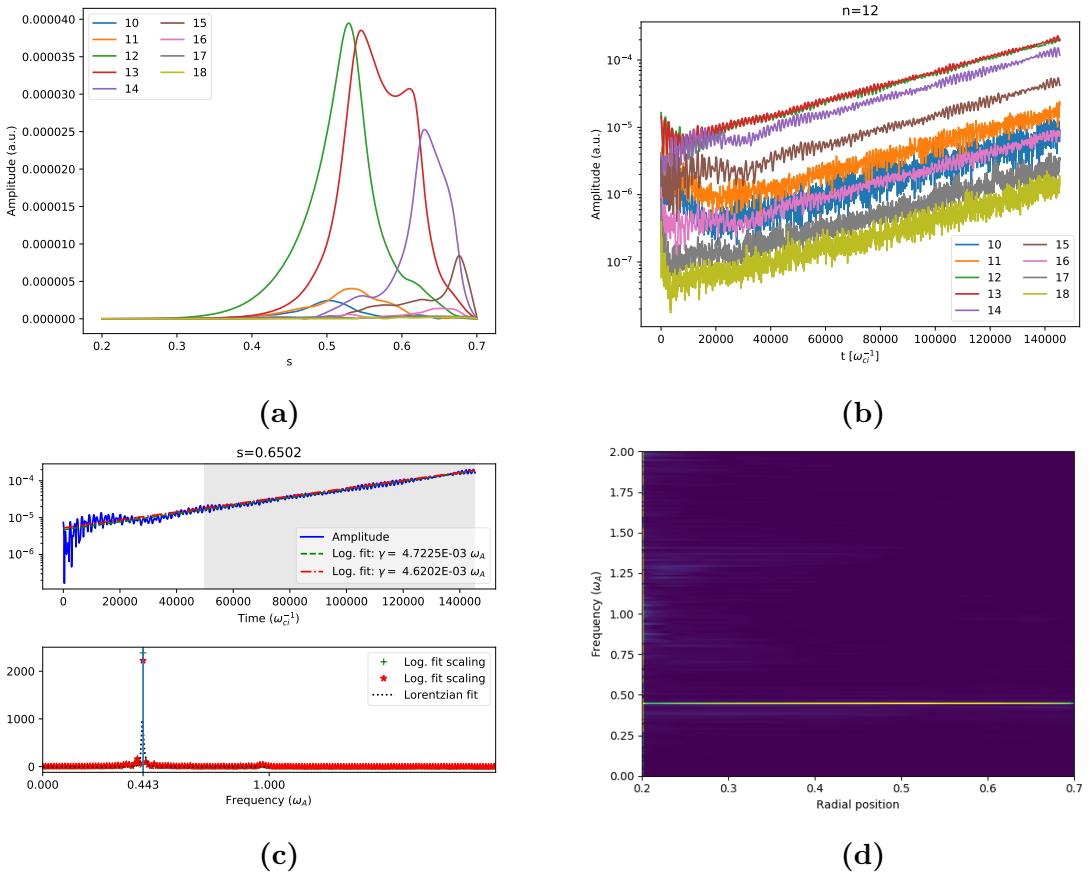

Figure 5.17: $n = 12$ TAE simulated in an annular domain ($0.2 \leq s \leq 0.7$). (a) radial structure of poloidal harmonics, (b) temporal evolution of peak values for each poloidal harmonic, (c) growth rate and frequency fit at mode position, (d) radial frequency spectrum of mode.

were performed but are omitted as the modes are very close to marginal stability, and the uncertainty on the fit is significant.

We also include on the figure some points showing the effect of finite Larmor radius (FLR) effects on the EP distribution. Here we see that for large n, this provides a significant reduction in the effectiveness of the mode drive. This shifts the peak in the mode growth rates to $n \approx 20 - 25$. This stabilization occurs as the gyro-radius approaches the scale of the mode structure. The offset between cases with zero Larmor radius (ZLR) and FLR is broadly inline with previous results [20], which reported a decrease in the linear drive of approximately 1% of the linear frequency for cases with nominal EP density for $20 \leq n \leq 30$. However, a quantitive comparison is left open for future work, as this should also consider non-perturbative effects.

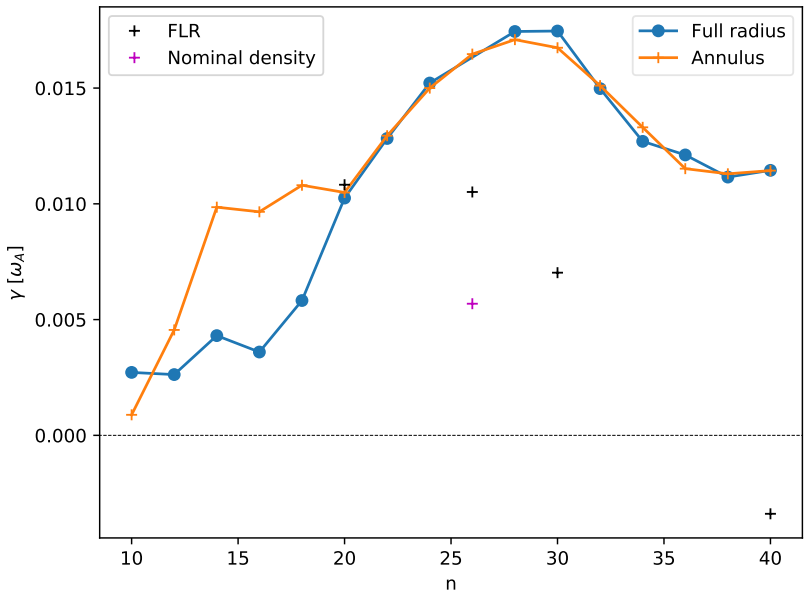

(a)

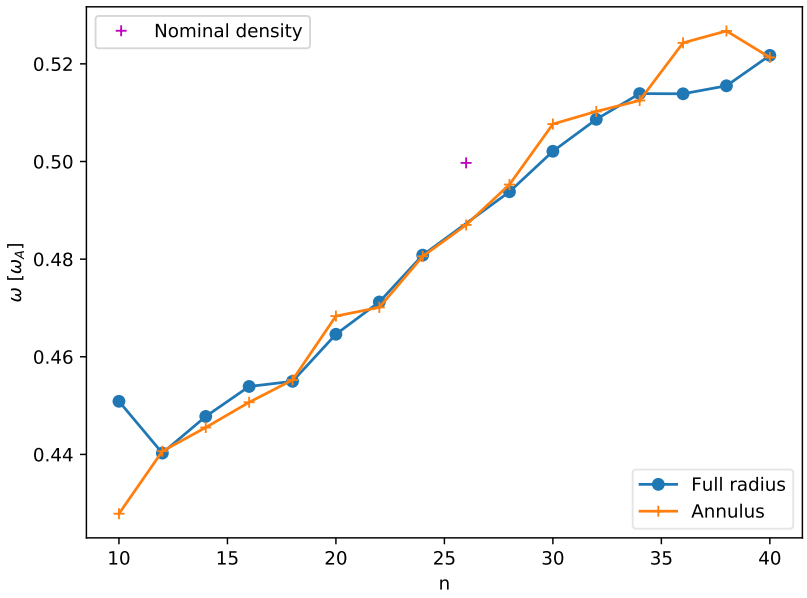

(b)

Figure 5.18: Scan of the mode growth rates (a) and frequencies (b) over a range of toroidal mode numbers n. These are calculated for simulations on both annular (crosses) and global (circles) domains. An additional point is added for reference, showing the results for $n = 26$ with nominal EP density on the annular domain (as in figure 5.11c). In (a), we also add growth rates measured with finite Larmor radius (FLR) effects enabled in the EPs (black crosses), simulated on annular domains.

6 Nonlinear global gyrokinetic simulations of energetic particle driven modes

Following on from the previous chapter, we can consider the nonlinear behaviour of the TAE spectrum shown in figure 5.18. Here, we run simulations with the wave-particle nonlinearity enabled, such that the energetic particles are free to become trapped in the perturbed fields, and therefore locally flatten the distribution function close to the resonances. This flattening removes the mode drive and therefore the mode saturates. This flattening also causes some amount of EP transport, and the local flattening can potentially create a localized steepening either side of the flattened region that can possibly lead to enhanced drive of the other weakly damped or linear subdominant modes.

6.1 Nonlinear single-mode simulations

The linear results discussed in §5.4 were, in fact, analyses of the linear phase of runs using the nonlinear equations of motion for the EPs in ORB5. It is therefore possible, by running the simulations longer, to observe mode saturation due to wave-particle nonlinearities. Given the results of §5.4, we focus initially on the most unstable modes, starting with mode numbers $20 \leq n \leq 30$.

In figure 6.1, we show the nonlinear evolution of the poloidal harmonics of the $n = 30$ case, initialized with $m = (30, 31)$, $m = (31, 32)$, $m = (32, 33)$ perturbations at radial positions $s \approx \{0.42, 0.55, 0.59\}$ respectively. Here, we see the subdominant TAEs, associated with $m = 32$–39, remain present, but saturate later at a lower amplitude than the TAE of the main gap, $m = (30, 31)$, although $m = 32$ has by this point become masked by the harmonic of the main TAE. We measure that the peak amplitudes for the $m = 30, 31$ harmonics are approximately $0.025, 0.0205$, with other harmonics approximately an order of magnitude or more smaller, with the largest harmonics which saturate later than the main gap, $m = 33$–36 with maximum amplitudes of approximately 0.001. Measuring the toroidal envelope, obtained by summing all the

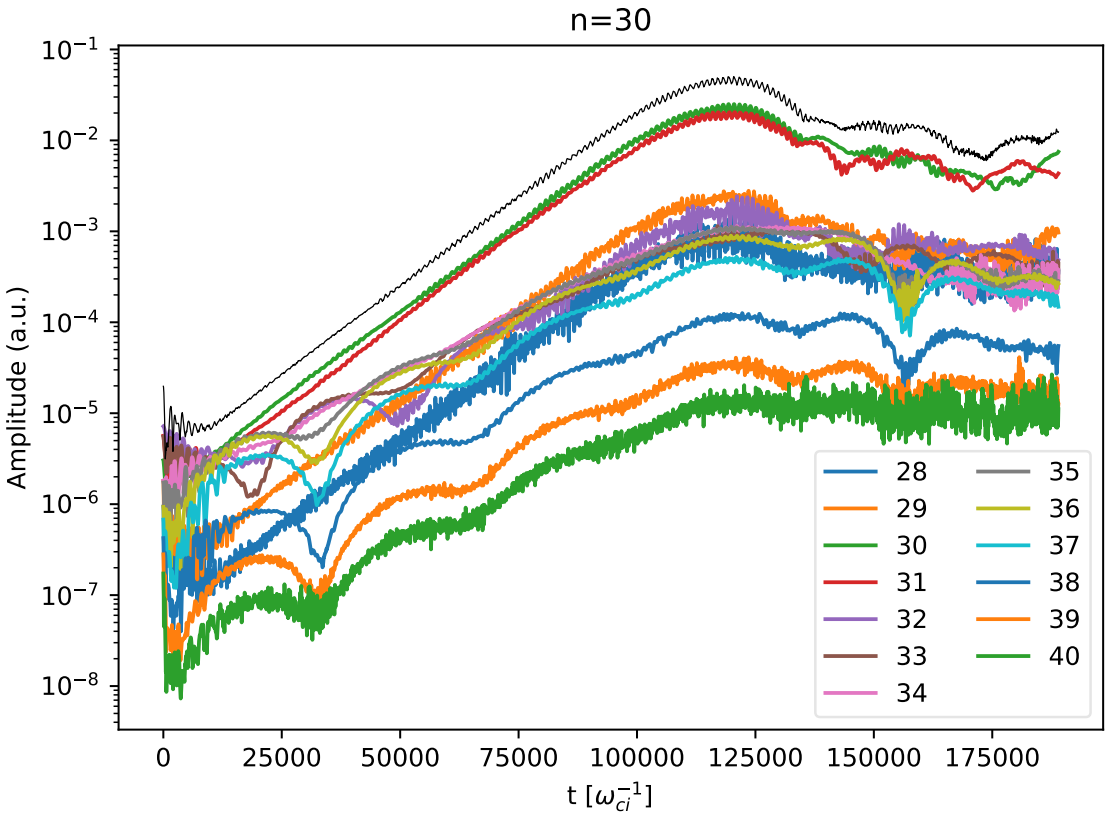

Figure 6.1: Nonlinear evolution of the poloidal harmonics of the electrostatic potential for $n = 30$ TAEs, run on an annular domain. In black, the peak value of the toroidal envelope is shown.

poloidal components for $n = 30$ on the outboard midplane of the low field side (LFS, poloidal angle $\chi = 0$), we observe a peak value of 0.05. These values of the saturation amplitude ($|\phi_n|$ or $|\phi_{mn}|$) are measured in arbitrary units, but in figure 6.2a, we see how this maps to the physical quantity of perturbed density $\delta n/n$, for which we see extrema of $-1.9 \cdot 10^{-3}$ and $2.5 \cdot 10^{-3}$.

In figure 6.2a, we show the initial and final profiles of the EP density, normalized to the volume averaged electron density. Close to the mode position ($s = 0.42$), we do not observe any significant change in the density profile caused by transport of EPs. Only a closer look at the perturbed density, shown in figure 6.2b, reveals that there is a small flattening in the EP density profile, as expected by the wave-particle saturation.

We notice non-perturbative features in the mode structure and frequency, noting the splitting of the main peaks of the harmonics in the frames of figure 6.3. From this pattern, we infer frequency chirping, the downward and upward frequency shifting of the inner and outer harmonic respectively.

We compare this result to the same case run on the full radial domain, noting a different initialization, here initializing only the TAE of the main gap, $m = (30, 31)$ at $s \approx 0.42$. When we look at the evolution of the poloidal harmonics of the potential, shown in figure 6.4, we see very similar behaviour. A quantitative comparison shows that peak values for $m = 30, 31$ are $0.025, 0.0204$, with $m = 29$ the next largest harmonic at approximately 0.003. In this case, the peak value of the toroidal envelope is also 0.05. This close match justifies the use of reduced radial domains for simulations of high n modes.

We also consider the case of $n = 20$, recalling from figure 5.18 that the global and annular results for $n = 20$ showed good linear agreement, whereas $n < 20$ showed significant discrepacies. We therefore compare the behaviour for annular and global $n = 20$ simulations in the nonlinear phase. This comparison is presented in figure 6.5. We see a reasonable agreement between the two cases, but note that some differences are can be seen.

6.2 Nonlinear multi-mode simulations

Having shown in the previous section reasonable agreement for the somewhat global TAE ($n = 20$), and excellent agreement for the well localized TAE ($n = 30$) between global and annular simulations, we justify the use of an annular domain for a multi-mode investigation into this range of mode numbers. Here, we run simulations with multiple modes present, keeping $n = \{20, 21, \ldots, 30\}$ in the filter when solving the field equations. We show the evolution of the peak envelope for each toroidal mode in figure 6.6. The values are calculated in the same way as the black line shown in figure 6.1. Here, we see

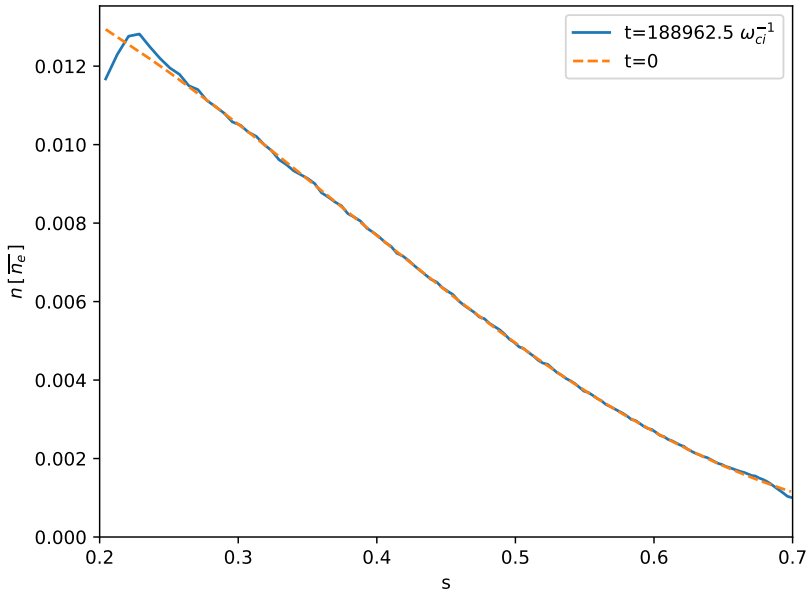

(a)

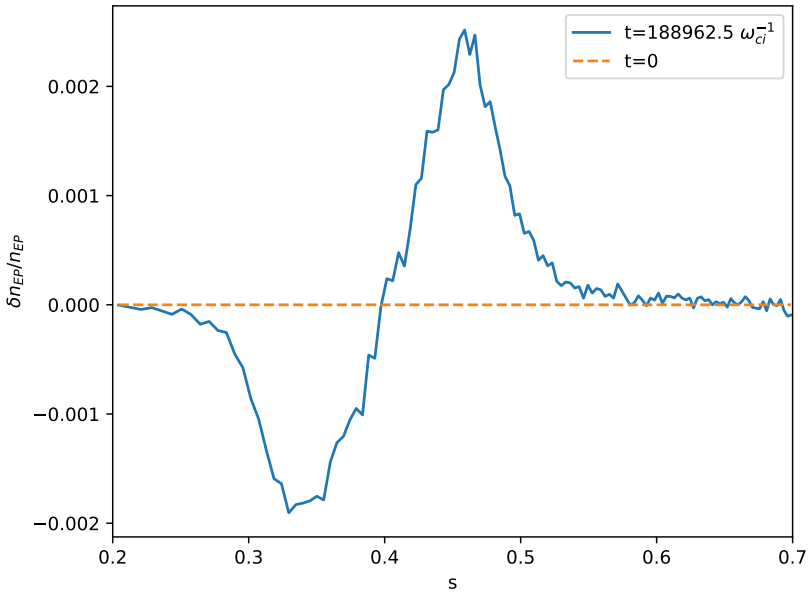

(b)

Figure 6.2: Initial and final profiles of the (a) total and (b) normalized perturbed EP density for the simulation shown in figure 6.1.

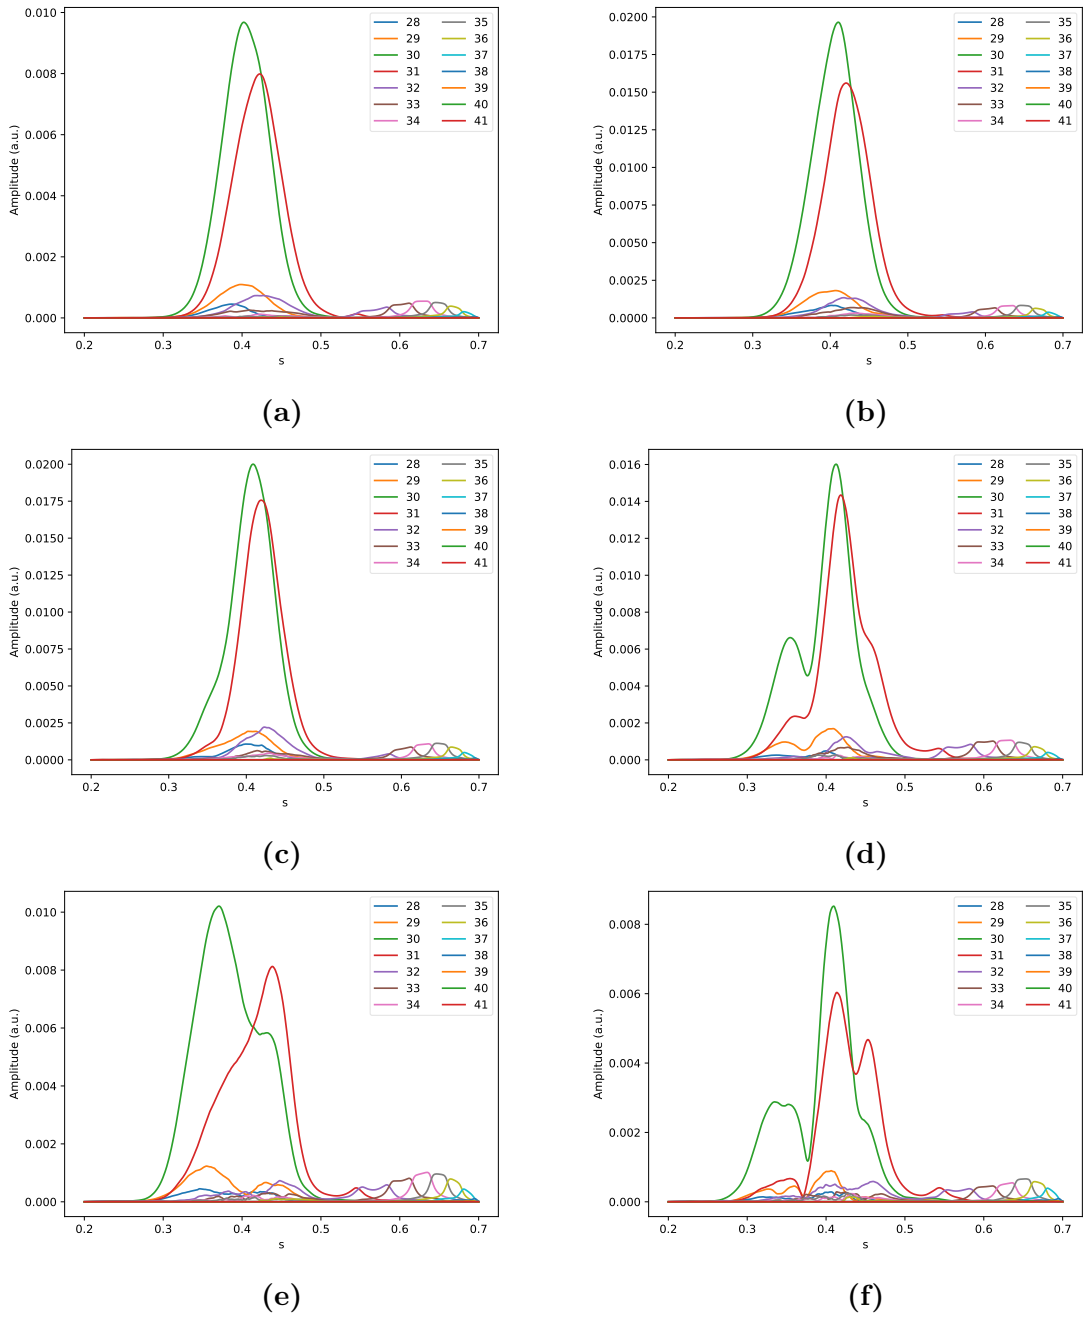

Figure 6.3: Snapshots of the radial structure of the poloidal harmonics of the electrostatic potential of the $n = 30$ case shown in figure 6.1, taken at times: (a) 100000, (b) 110000, (c) 120000, (d) 130000, (e) 140000, (f) 150000 ω_{ci}^{-1}.

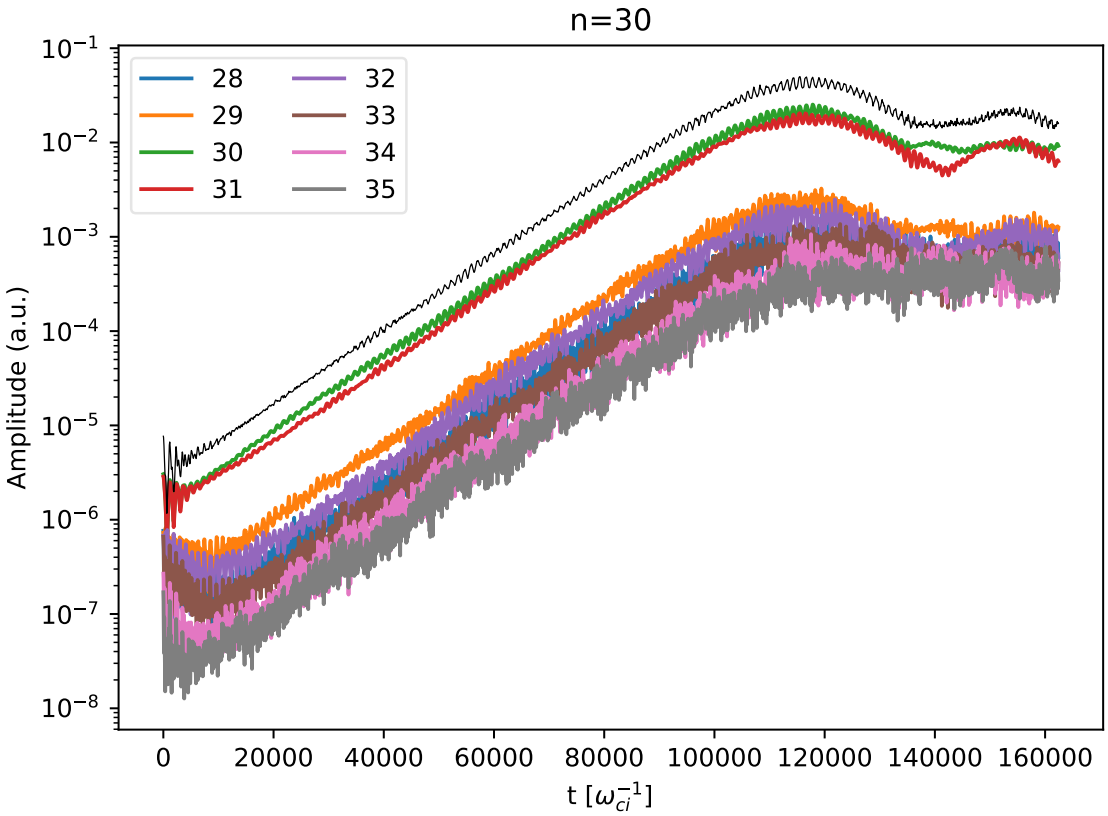

Figure 6.4: Nonlinear evolution of the poloidal harmonics of the electrostatic potential for $n = 30$ TAEs, run on an full radial domain. In black, the peak value of the toroidal envelope is shown.

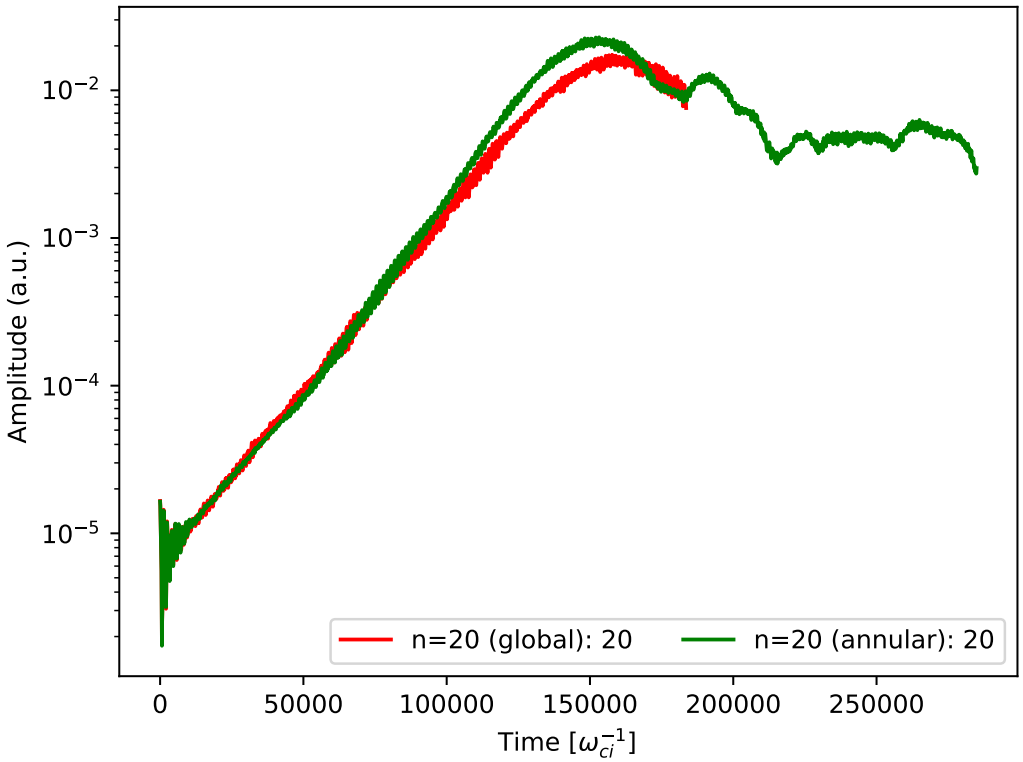

Figure 6.5: Nonlinear evolution of the toroidal envelopes (LFS) of the electrostatic potential for $n = 20$ TAEs, simulated on full radius and annular domains.

that the modes initially saturate at values of $\approx 2\text{–}5 \cdot 10^{-2}$, before undergoing another growth phase, and reaching values almost an order of magnitude larger, noting that $n = 20$ and $n = 23$ appear to show the largest values.

To investigate this, we look at the evolution of the poloidal harmonics of some of the modes. In the upper panels of figure 6.7, we show the evolution of the $n = 20$ and $n = 24$ components of the potential from this multi-mode simulation, plotting in each case various poloidal harmonics. In the lower panels, we plot the same, but for the cases with $n = 20$ or $n = 24$ only in the simulation, in these cases also overlaying the peak of the toroidal envelope in black. In the nonlinear phase there are some clear differences between top and bottom figures. In the linear phases of all the simulations, we see the potential is dominated by the $m = (n, n + 1)$ main gap TAE. In the single-mode simulations, this mode saturates and dominates, just as we saw in figure 6.1 for $n = 30$. These saturate at comparable amplitudes in the single and multi-mode cases. However, in the upper panels, we see that the higher harmonics, belonging to the outer TAE gap(s) in the system continue to grow after the initial saturation of the main gap TAE. These modes are then responsible for the overall increase in the envelopes, as shown in figure 6.6.

These features are confirmed in the panels of figure 6.10, where for $n = 20$ and $n = 24$, we plot the radial structure of the poloidal harmonics during the linear phase, close to peak saturation, and in the late time. In the linear phase (upper panels, $t \approx 80000\,\omega_{\text{ci}}^{-1}$), where we see the dominant peak is due to the $m = (n, n + 1)$ harmonics of the main TAE gap at $s \approx 0.49$ and $s \approx 0.47$ for $n = 20$ and $n = 24$ respectively. In black we overlay the toroidal envelope on the low field side (LFS, poloidal angle, $\chi = 0$) in solid, and on the high field side (HFS, $\chi = \pi$) in dashed. We observe, as before, that the even parity TAE, peaked on the LFS is dominant. Close to the peak saturation (middle panels, $t \approx 200000\,\omega_{\text{ci}}^{-1}$), we see that for both $n = 20$ and $n = 24$, we observe the peak further outside, $s \gtrsim 0.55$, on the LFS. These are the perturbations which exhibit enhanced saturation in figure 6.7. In the late time (bottom panels, $t \approx 390000\,\omega_{\text{ci}}^{-1}$), we see that the outside peaks are still relatively large, although smaller than in the middle panels, but we also see that in the main gap, the harmonics $m = (n, n + 1)$ exhibit odd parity, peaked on the HFS. This observation reminds us of the behaviour shown in §5.2.2, where we observe that, even though the odd parity TAE is not effectively driven linearly by the EPs, it is weakly damped.

We also plot in figure 6.9 the contributions to the electrostatic potential in the poloidal plane of $n = 20, 24, 30$, and the total poloidal plane. In particular, in the total poloidal plane image, figure 6.9e, we observe triangular, or boomerang-like features [103] in the mode structure, associated with the mode symmetry breaking due to radial propagation caused by a difference in the radial position of mode drive and damping [104].

We perform a type of frequency analysis, different here from the previous frequency

analyses because of both the multi-mode simulation nature and the mixed LFS/HFS peaking at different times. In this analysis, shown in figure 6.11, we plot the evolution of the instantaneous frequency, as measured by evaluating the time derivative of the phase of the Fourier component of the potential. This method can give an accurate measure of the frequency if the signal contains only a single frequency. This evaluation is performed at $s = 0.489$, the peak of the linear mode structure. Here we clearly observe a change in the frequency during the late nonlinear time, with an emerging dominant frequency at late time of $0.64\,\omega_{\mathrm{A}}$, consistent with the odd TAE sitting at the top of the gap in the continuum.

As we see in figures 6.9, 6.10, 6.11, and 6.12, there are clear signs of non-perturbative modifications to the mode structure and the frequency. Such features cannot be described as a superposition of $m = (n, n+1)$, $m = (n+1, n+2)$, ... gap modes. A key observation is that such non-perturbative behaviour is observed in cases relatively close to marginal stability ($\gamma_{lin,max}/\omega_{lin} \approx 3\%$).

The energetic particle transport, i.e. the modification to the EP density profile, is shown in figure 6.8. We see a significant amount of redistribution in the region $0.5 \lesssim s \leq 0.7$, and evidence of significant losses (relative to our domain) at the outer boundary ($s_{\max} = 0.7$). The maximum change in the EP density is observed at $s \approx 0.55$, where we observe $\delta n_{EP}/n_{EP} = -0.34$.

Although the single-mode nonlinear saturation for $n = 30$ matched closely between annular and global domain simulations, and that the linear growth rates matched well for $n = 20$–30; in the multi-mode case, given the increased transport at the edge of the domain, non-local to the linearly most unstable modes, it is an open question as to how these results would be modified in the case of considering a global simulation domain. The multi-mode simulation presented in this section was particular numerically expensive, being run with $10\times$ the number of markers as the single-mode cases $N_{\mathrm{p}} = \{320, 1280, 320\} \cdot 10^6$ for the main ion, electron and energetic particle species respectively. This makes it very difficult to perform a detailed sensitivity study, in particular ruling out global runs at this time.

We do, however, attempt to evaluate sensitivity with respect to the retained modes in the system by simulating the same case, but keeping only every second mode ($n = \{20, 22, \ldots, 30\}$), which we run with half the number of markers as the previous case. The evolution of the toroidal envelopes is shown in figure 6.13a. Here, we see qualitatively similar behaviour. We compare also the EP transport, plotting in figure 6.13b in blue the EP density profile at the end of the simulation ($t \approx 200000\,\omega_{\mathrm{ci}}$), which we compare to the thin black line, which shows the EP density of the case in figure 6.8 as measured at the same time as the blue line. Here we see that the redistribution matches almost exactly in the case of with only even n as compared to that when we keep all n. This increase redistribution is observed to be a robust feature

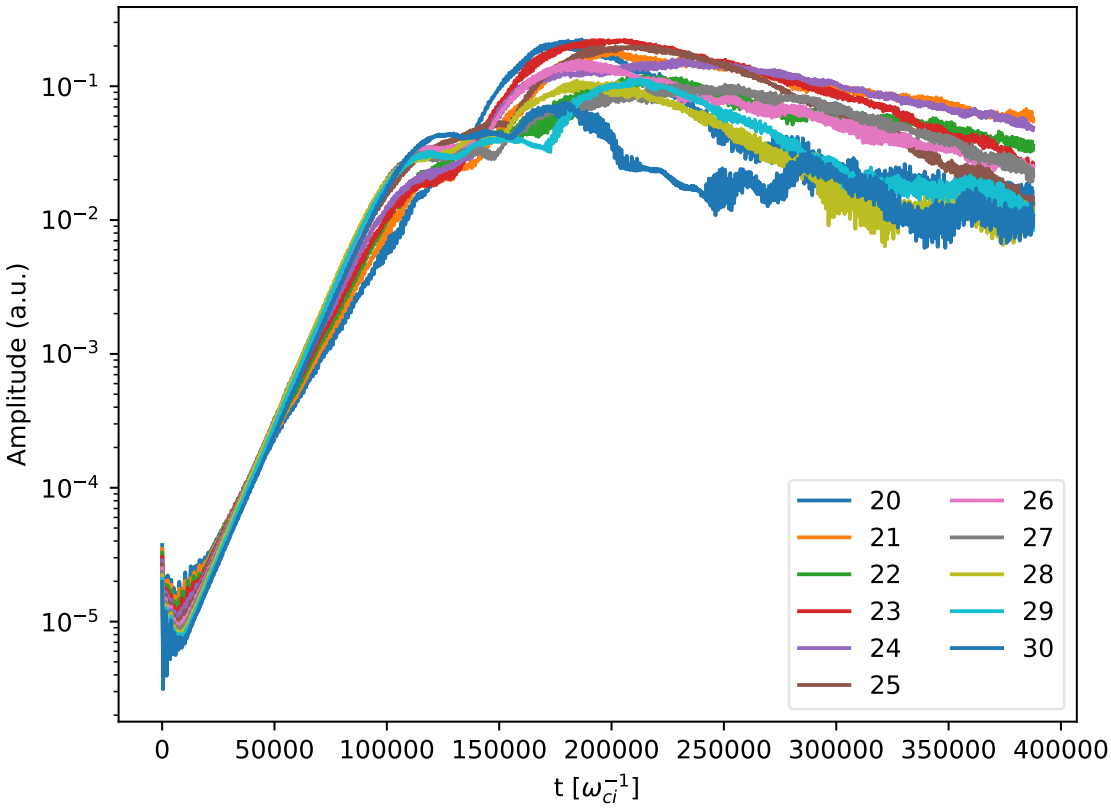

Figure 6.6: Nonlinear evolution of the peak values of the toroidal envelope of each toroidal mode n, simulated on an annular domain.

of multi-mode nonlinear simulations.

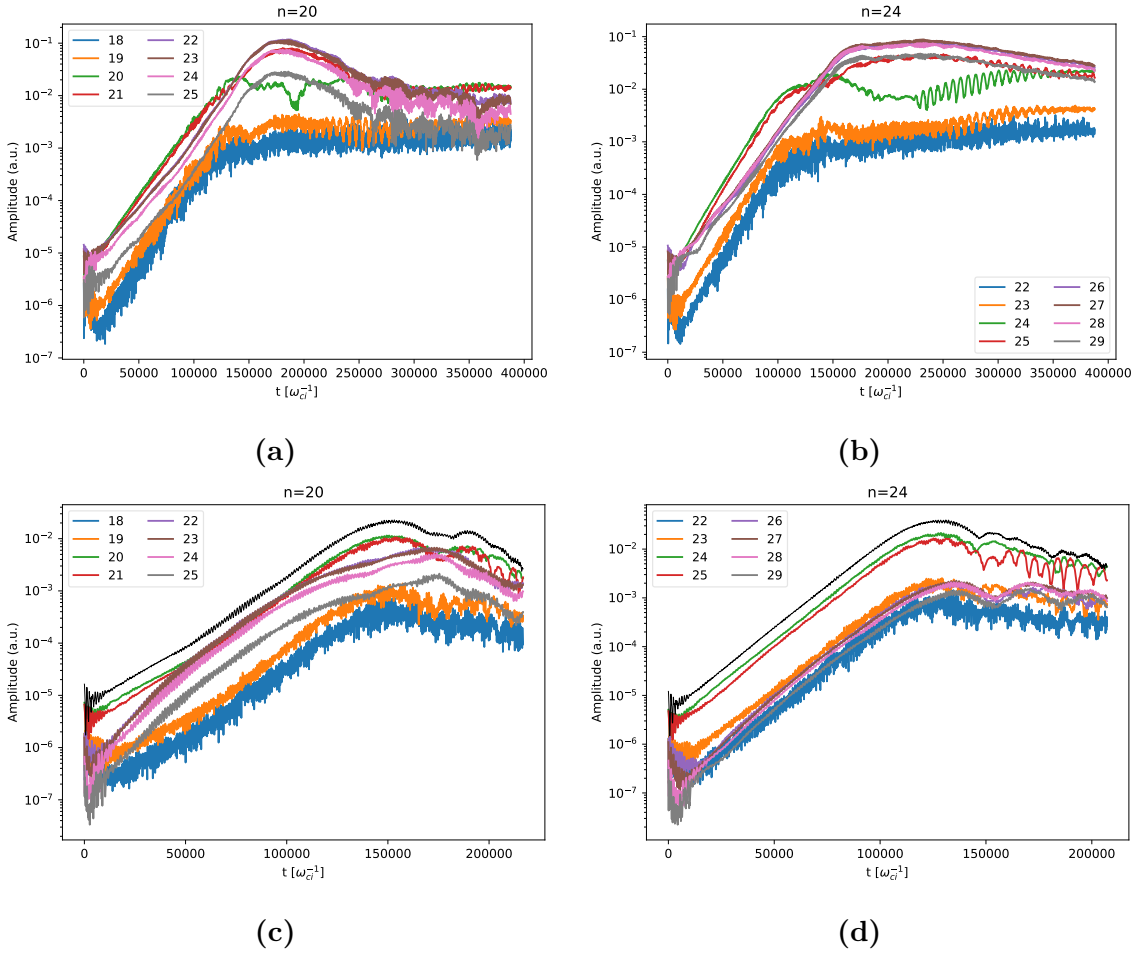

Figure 6.7: Nonlinear evolution of the poloidal harmonics of (a) $n = 20$ and (b) $n = 24$ from the $n = \{20 \dots 30\}$ simulation shown in figure 6.6. (c) and (d) show the nonlinear evolution of the poloidal harmonics of $n = 20$ and $n = 24$ respectively in single mode simulations. In black, the toroidal peak is plotted.

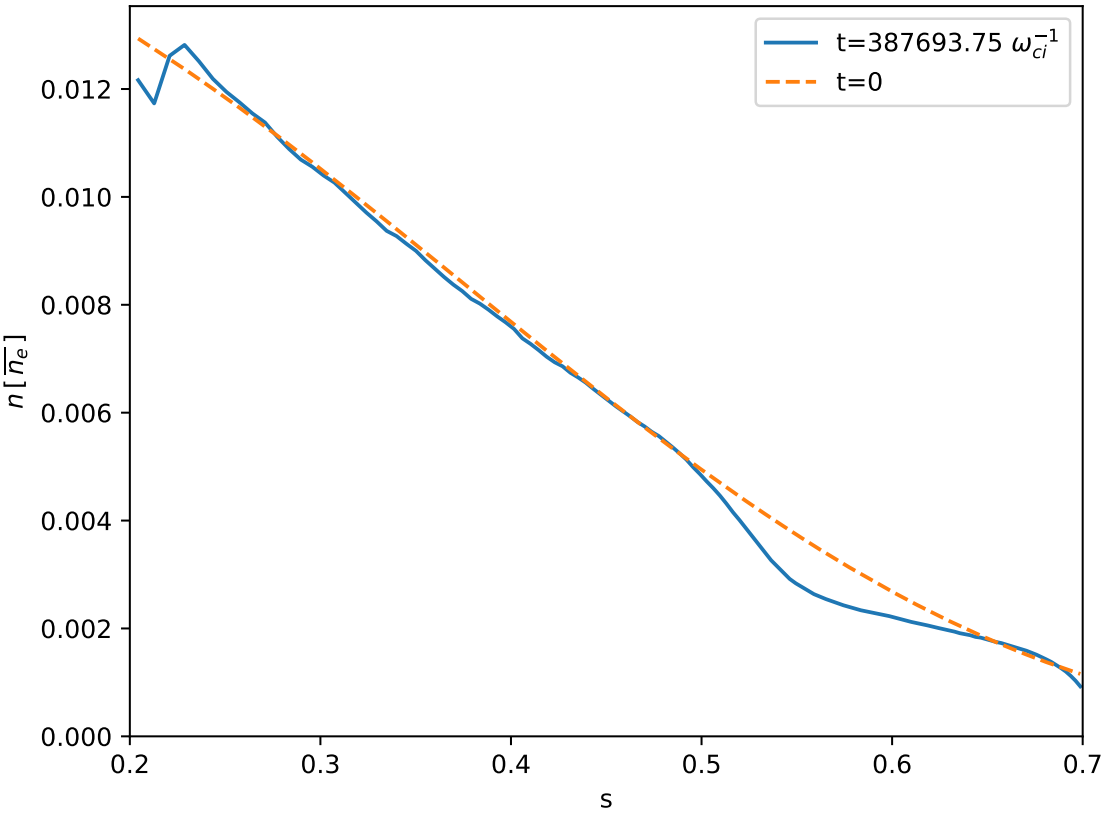

Figure 6.8: Initial and final profiles of the EP density for the simulation shown in figure 6.6.

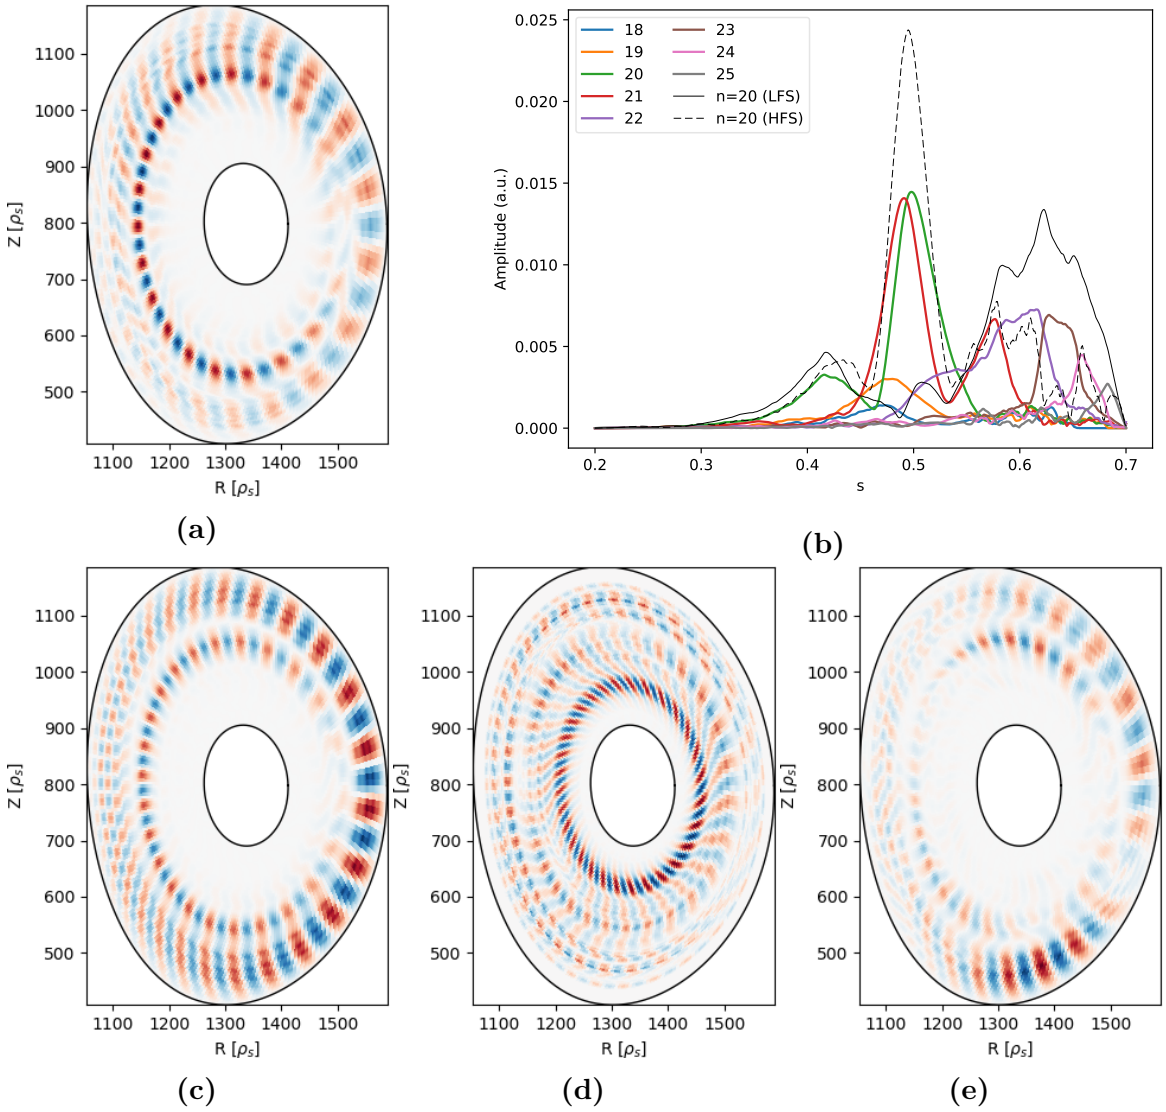

Figure 6.9: Contributions to the electrostatic potential at the late nonlinear phase ($t = 387675\,\omega_{ci}^{-1}$). (a) shows the poloidal plane, showing $n = 20$, and (b) shows the poloidal harmonic structure, also for $n = 20$, overlaid with a solid black line showing the toroidal envelope at the low field side (LFS) (as was done in e.g. figure 6.1), and with a dashed black line showing the toroidal envelope caluclated at the high field side (HFS), (poloidal angle $\chi = \pi$). (c) shows the $n = 24$ contribution to the poloidal plane, (d) the $n = 30$, and (e) the contributions from all n. These data are taken from the $n = \{20 \ldots 30\}$ simulation shown in figure 6.6.

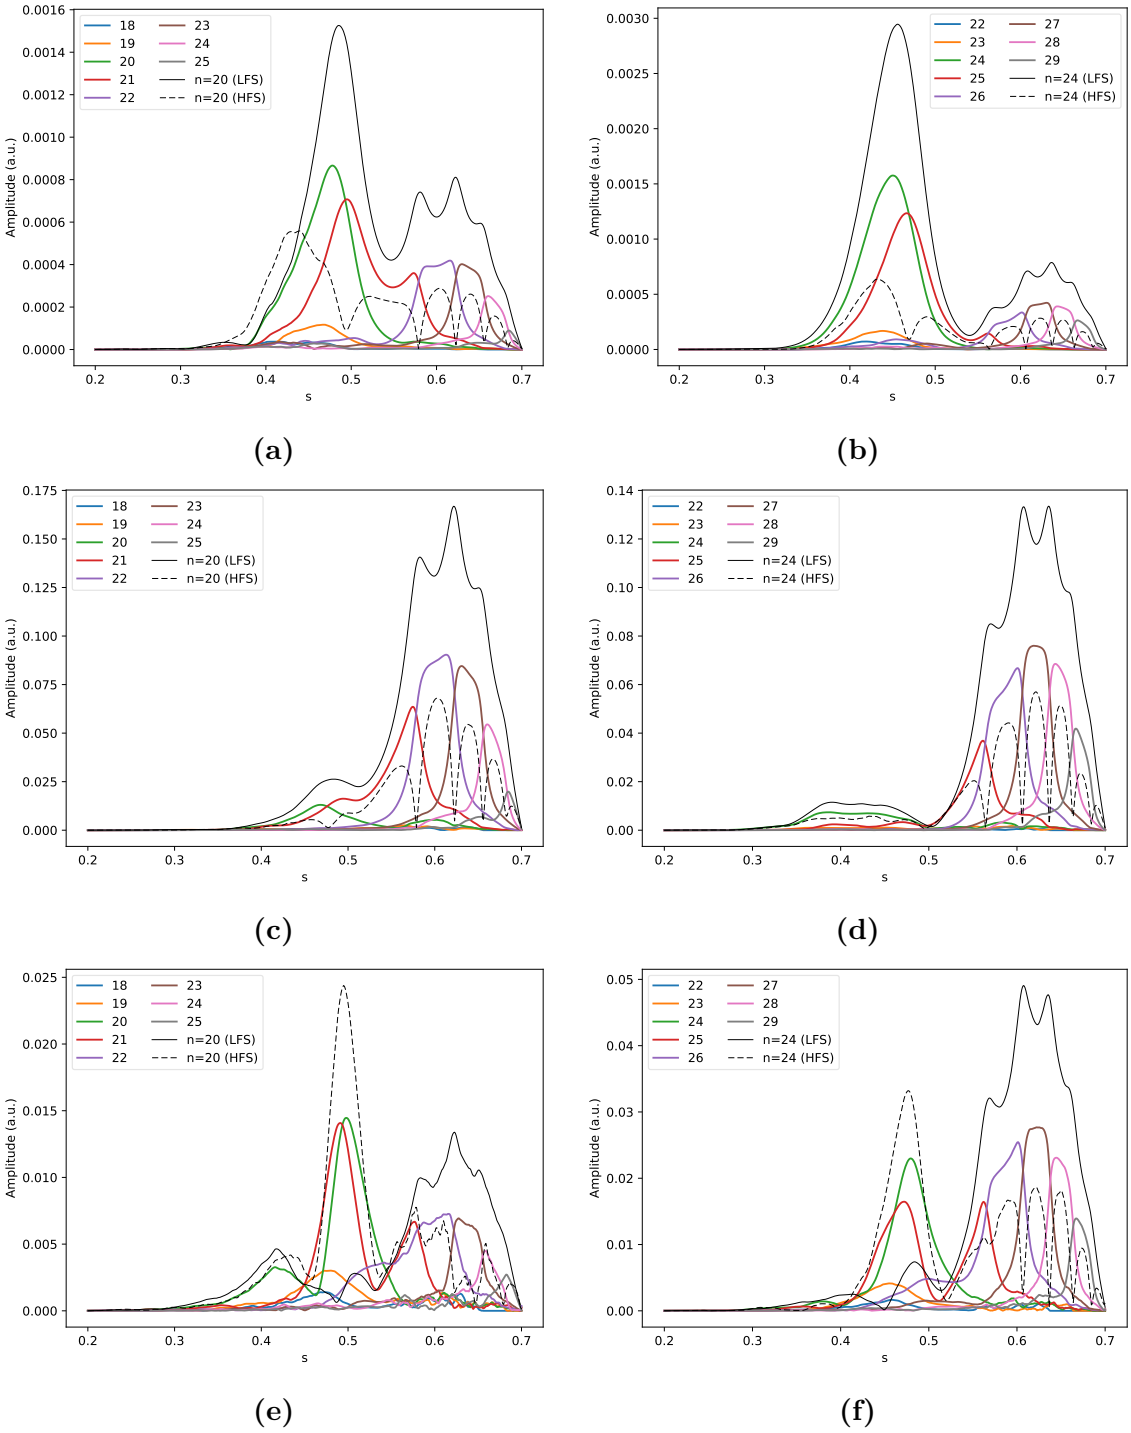

Figure 6.10: Poloidal harmonics of the electrostatic potential for (left) $n = 20$ and (right) $n = 24$. These are plotted at (top) $t = 79987.5\,\omega_{\mathrm{ci}}^{-1}$, (middle) $t = 199987.5\,\omega_{\mathrm{ci}}{}^{-1}$, and (bottom) $t = 387675.0\,\omega_{\mathrm{ci}}^{-1}$. In solid black, we overlay the LFS toroidal envelope at $\chi = 0$, and in dashed black the HFS toroidal envelope ($\chi = \pi$). These data correspond to the simulation in figure 6.6.

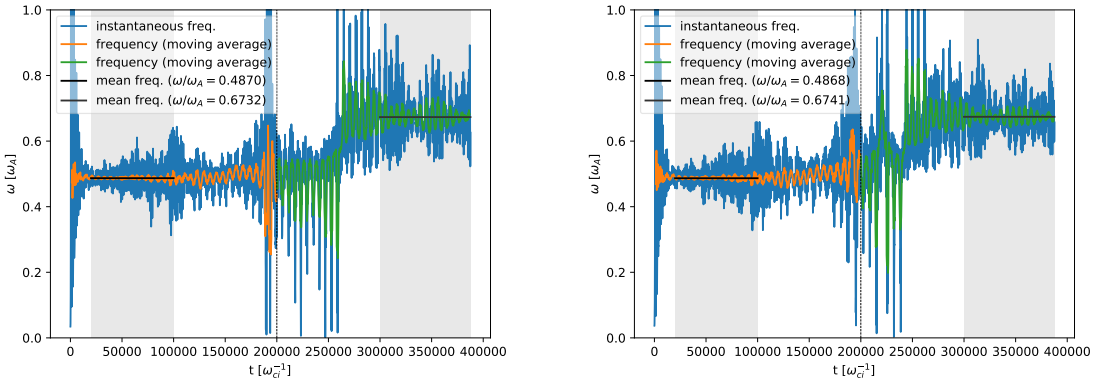

(a) **(b)**

Figure 6.11: Instantaneous frequency, measured as the absolute magnitude of the time derivative of the phase of the Fourier coefficient of the potential for $n = 20$, calculated using (a) $m = 20$ and (b) $m = 21$. The frequencies are fit in the shaded regions, shown in black and grey lines for the linear and late nonlinear regimes respectively. The orange and green lines show the moving average of the instanteous frequency averaged over $T = 2\pi/\omega_{\mathrm{av}}$ for $\omega_{\mathrm{av}} = 0.463$ and $0.64\ \omega_{\mathrm{A}}^{-1}$ respectively, changing at $t = 200000\,\omega_{\mathrm{ci}}$ (the dashed vertical line). Data from the simulation shown in figure 6.6.

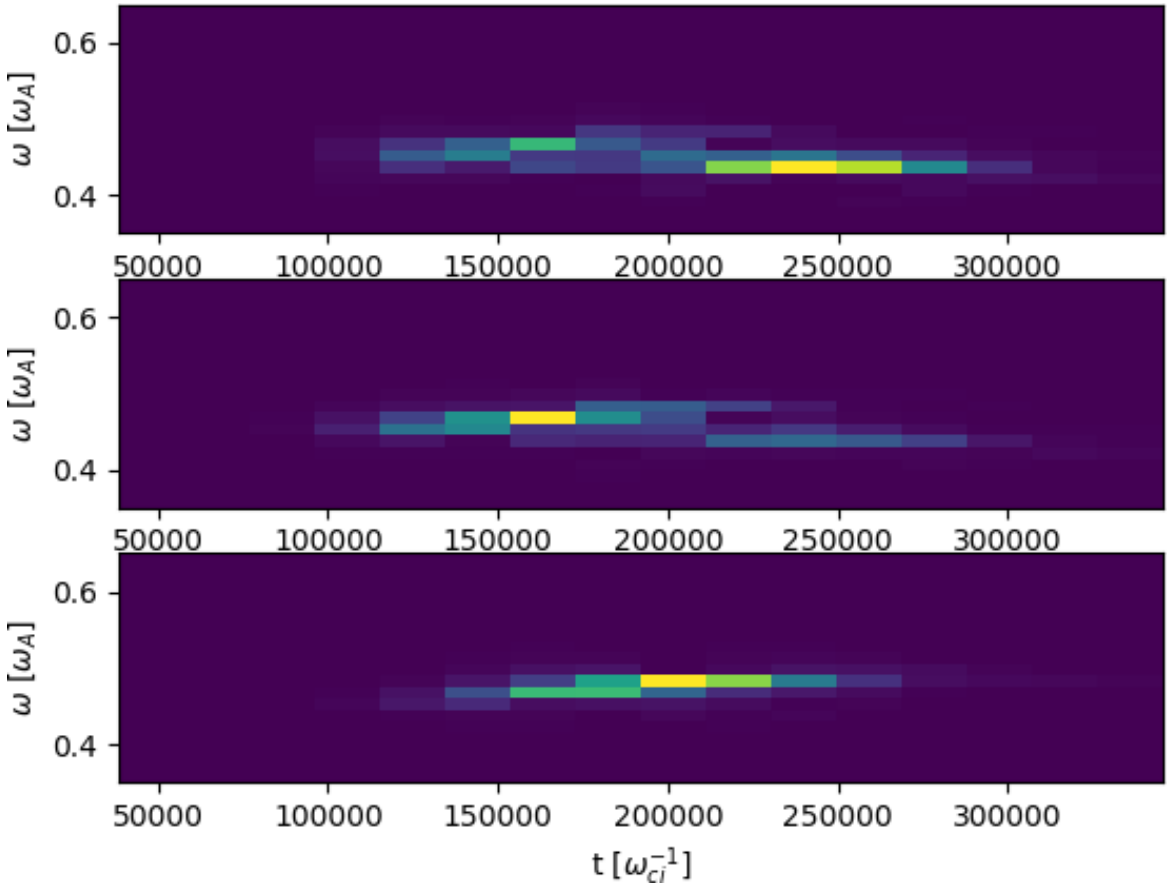

Figure 6.12: Short time Fourier transform (STFT) of the $n = 20$ signal on the LFS measured at $s = 0.45$, 0.489, and 0.53 (top to bottom). Data from the simulation shown in figure 6.6.

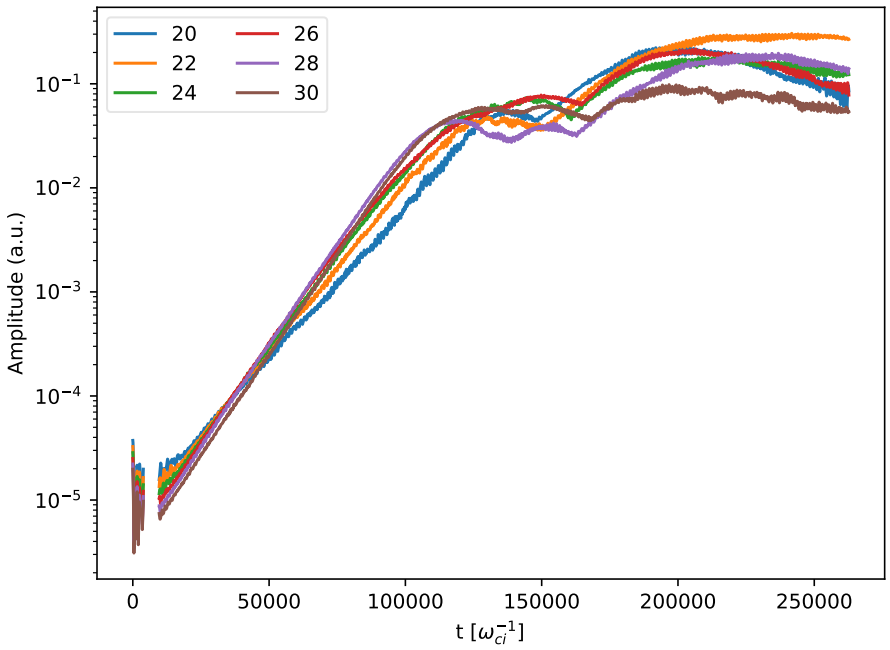

(a)

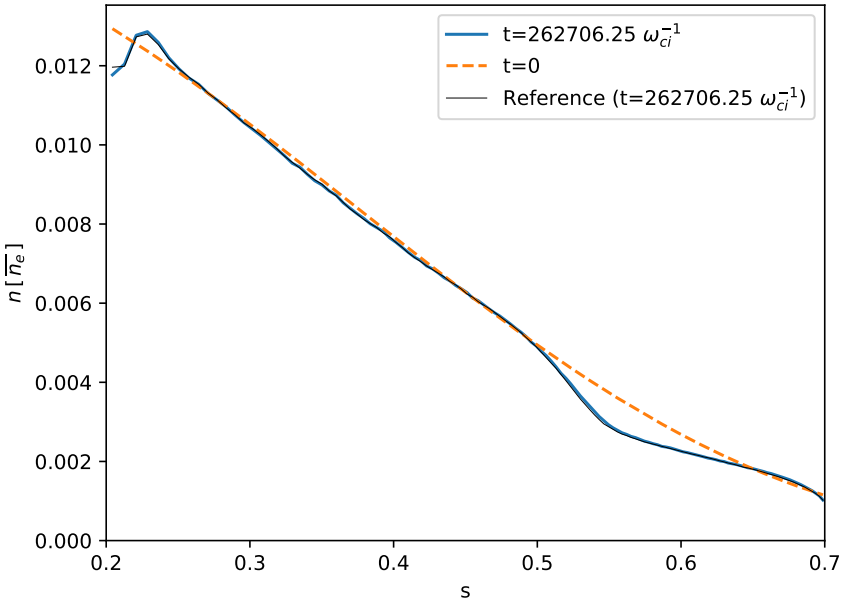

(b)

Figure 6.13: Time evolution of toroidal envelopes (a) and the initial and final EP density profiles (b) for a nonlinear annular simulation with $n = \{20, 22, \ldots, 30\}$. In (b), we overlay the EP density profile from the case with $n = \{20, 21, \ldots, 30\}$ (as shown in figure 6.6), as measured at the same time.

7 Discussion and summary

During this thesis, we have shown a numerical study of the impact of TAE modes in the 15MA scenario for ITER, with double the nominal scenario value for the EP density.

This has been performed in chapter 4 using various flavours of a hybrid model, based on the linear gyrokinetic eigenvalue solver LIGKA and the nonlinear perturbative wave-particle hybrid initial value code HAGIS. Considerable effort was put into automation of the hybrid model, as well as evaluating the sensitivity with respect to the model used for the eigenvalues, with view to performing automatic evaluation of plasma scenarios, or even embedding into a future EP transport model. We find that for well localized gap modes (in the considered ITER scenario, this means $n > 20$), the linear behaviour is not particularly sensitive to the model, with the global kinetic eigenfunctions delivering similar behaviour to those calculated in the MHD limit combined with local kinetic damping rates. We found that the TAE gaps in which modes are driven unstable depend on the equilibrium, with a flat q scenario showing also modes in the 'minus-oneth' gap and the 'first' gap in addition to the 'main', 'zeroth' TAE gap. For the main scenario, we see the unstable linear spectrum is dominated by the even parity modes of the main gap, although other branches may be less damped. For intermediate n ($10 \leq n \leq 15$), the growth rates for multiple TAEs gaps may be comparable. These results are similar to previous works, but we developed automated workflows, which allowed a more systematic evaluation of the results.

For the modes of the main TAE gap of the regular scenario, we performed nonlinear single-n simulations using the different models and compared the saturation amplitudes. Comparing the saturation amplitudes between global kinetic and MHD workflows, we observed a slightly higher saturation amplitude for intermediate n, $20 \leq n \leq 35$, with both workflows peaking around $n \approx 25$. In order to be able to easily predict the saturation amplitude, it would be convenient if possible to ascertain the scaling between linear properties such as growth rate or frequency and nonlinear saturation amplitude. Some analytical work predicts a scaling of $A_{\max} = C(\gamma/\omega)^2$ (depending on saturation mechanism), however these works typically look at the effect of changing the linear drive or damping of a single mode, where resonance properties will not change, rather than the scaling of many different modes with respect to their individual growth rates. Therefore deviations from this scaling are to be expected. We find that even under these assumptions, the scaling coefficient is constant (MHD workflow) or varies

smoothly (kinetic workflow). A difference of ≈ factor 2 between the coefficients with the two workflows is noted, and further work could investigate this discrepancy. However, due to the smoothness, it would be possible to simulate only a few cases nonlinearly and interpolate the coefficient between them to predict the saturation of all cases. This fast predictive ability could therefore be combined with a future transport model.

Finally we performed a numerical test designed to mimic the effect of mode structure modification as a result of non-perturbative nonlinear behaviour. We implemented the mechanisms for handling time dependent mode structures, and evaluate this by taking an analytical mode structure (Gaussian), with a smoothed step function for the width term. Performing this procedure with many different parameters (amount, time, and speed of the transition), we found that the peak saturation amplitude would follow the amplitude of a static case with the same width as the width at the time of the peak. This gives us confidence that we would be able to mimic non-perturbative nonlinear behaviour, even with a perturbative model. It also demonstrates the mechanisms by which we could update the mode structure as the simulation evolves leading to a more self-consistent coupled non-perturbative HAGIS/LIGKA model.

In chapter 5, we started evaluating the possibility of running the ITER scenario with a global gyrokinetic initial value code. An important step of the process was the extensive convergence tests performed, necessary for faith in numerical predictions, in particular given the extremely large numerical cost of running electromagnetic simulations with kinetic electrons. To begin, we made a reduction of the problem (Maxwellian EP distribution function, neglecting gyroaverage on EPs, hybrid background isotope, flattened temperature and density profiles for the background plasma, reduced mass ratio). We demonstated that we can observe AEs with correct linear physics properties, including higher frequency modes such as EAEs in decaying plasmas. We demonstated also subdominant modes, the odd-parity TAE, and agreement with the MHD/kinetic continuum spectrum. Moving then towards more realistic cases, we restored the temperature and density profiles (but removed the gradients in the steep edge region) and assessed the impact of the mass ratio on linear physics and numerical properties. Physically, we found a slight increase in electron Landau damping for cases with heavy electrons, and numerically we found the interesting result that the convergence with respect to the time step scales $\approx (m_e/m_i)$.

We performed a systematic linear scan of the toroidal mode numbers for both a full radial domain as well as running on a reduced annulus. For low mode numbers such as $n = 12$, the coupling of the outer poloidal harmonics leads to two or more similar eigenfunctions with slightly different mode structures, frequencies, but very similar growth rates. This phenomenon was observed clearly when running on the full radial domain, but not when the outer harmonics are cut when running on an annulus. For higher mode numbers ($n \geq 20$) growth rates and frequencies agree. We add for

comparison a case with nominal EP density, and some cases with the gyroaverage applied for the EPs (FLR effects). FLR effects reduce the linear drive for $n > 20$, such that $n = 40$ is fully stabilized. These linear studies demonstrate that we are able to reproduce the linear properties, thus validating the code to linear eigenvalue calculations. This gives us confidence in our predictive abilities when performing simulations of TAEs with ORB5.

In chapter 6, we performed nonlinear continuations of the annular cases from the previous chapter. As an example, $n = 30$ was run to nonlinear saturation, and we observe that the EP redistribution is small. We compared this case with a nonlinear full radius case, with complete agreement. We observed clear indications of non-perturbative features in the mode structure and the frequency.

We then showed the impact of nonlinear interaction of multiple TAEs, running a nonlinear multimode simulation with $n = 20, \ldots, 30$. We saw a significant increase in the total amplitude, with a 'second saturation phase' not seen in single-n simulations. These correspond to an enhancement of the drive of the outer gap TAEs caused by redistribution from the linearly dominant TAEs of the main gap. In the late (decaying) phase, we observed the presence of odd-parity TAEs, reminding us of the previous observation in decay experiments. In this case, we observed a significant redistribution of EP density.

Finally, a similar simulation was performed, but keeping only every second mode. The total redistribution was very similar to the dense case.

These nonlinear simulations with ORB5 show some agreement with previous nonlinear hybrid works, for example the transition to regimes where we observe nonlinear excitation of linearly subdominant modes leading to an overall enhancement in the total EP transport, but also demonstrate features not observed in perturbative hybrid modelling, such as the mode structure perturbation and frequency chirping.

We stress that we do not predict any danger for ITER with nominal parameters as the findings showing enhanced nonlinear redistribution were evaluated at doubled EP density. Instead, we demonstrate that it is possible to observe an enhancement due to multi-mode effects, in contrast to some previous claims that such a threshhold may be more than an order of magnitude larger [18].

To the best of our knowledge, no global gyrokinetic initial value code has previously been applied to the TAE problem in ITER, either linearly or nonlinearly. This work therefore goes beyond linear eigenvalue, linear and nonlinear hybrid-kinetic, and linear local gyrokinetic treatments applied to this problem in the past and therefore offers a quantitative improvement in our abilities to predict the transition between different EP transport regimes.

7.1 Outlook

Unfortunately, this work is not the end of the story. It does, however, provide the foundation upon which to address the open issues which still remain.

As was mentioned in numerical introduction of the ORB5 code, small toroidal mode numbers (especially $n = 0$) are particularly difficult to handle in global electromagnetic simulations with large β. In this work, we have been able to run cases with $n = 6$, and even then only after some improvements had been made.

The nonlinear wave-wave coupling effects require the retention of small n in order for three-wave coupling to be evaluated for similar n, and $n = 0$ in order to include the self-coupling. Therefore the impact of wave-wave nonlinearities was not investigated in this work. In cases with small plasma β, studies including the effect of $n = 0$ recently been performed using the ORB5 code.

We believe that the issues related to the handling of $n = 0$ with large β in a global EM model will require a more careful handling of the equilibrium and the profiles, and the distribution functions must be exactly consistent with the equilibrium.

Another limitation of our work was the restriction to Maxwellian distribution functions. Although this is not a fundamental limitation of the model or the code, quantitative predictions for ITER TAE stability should follow the implementation of a more realistic distribution function, although the comparisons of the mode drive calculated using HAGIS provides a context for our quantative study.

While we showed results of multi-mode simulations for the modes $n = 20 \ldots 30$, other choices could have been made. Hybrid multi-mode simulations have previously shown the influence of the choice of included modes when performing multi-mode nonlinear studies.

For longer time simulations, the lack of sources and collisions begin to play a role. These effects could be included, although for the time scales included in this work, their effects are small.

Bibliography

[1] ITER Physics Basis Editors, ITER Physics Expert Group Chairs and Co-Chairs, and ITER Joint Central Team and Physics Unit. "Chapter 1: Overview and summary", *Nuclear Fusion* 39, 12 (1999).

[2] A. Gibson et al. "Deuterium–tritium plasmas in the Joint European Torus (JET): Behavior and implications", *Physics of Plasmas* 5, 5 (1998).

[3] R. J. Hawryluk et al. "Fusion plasma experiments on TFTR: A 20 year retrospective", *Physics of Plasmas* 5, 5 (1998).

[4] A. Polevoi et al. "ITER Confinement and Stability Modelling", *Journal of Plasma and Fusion Research Series* 5, (2002).

[5] G. Pereverzev et al. "Theoretical predictions of the density profile in a tokamak reactor", *Nuclear Fusion* 45, 4 (2005).

[6] C. Angioni et al. "Particle pinch and collisionality in gyrokinetic simulations of tokamak plasma turbulence", *Physics of Plasmas* 16, 6 (2009).

[7] E. Fable et al. "The role of the source versus the collisionality in predicting a reactor density profile as observed on ASDEX Upgrade discharges", *Nuclear Fusion* 59, 7 (2019).

[8] P. Lauber. "Local and global kinetic stability analysis of Alfvén eigenmodes in the 15 MA ITER scenario", *Plasma Physics and Controlled Fusion* 57, 5 (2015).

[9] M. García-Muñoz et al. "Scintillator based detector for fast-ion losses induced by magnetohydrodynamic instabilities in the ASDEX Upgrade tokamak", *Review of Scientific Instruments* 80, 5 (2009).

[10] C. S. Collins et al. "Observation of Critical-Gradient Behavior in Alfvén-Eigenmode-Induced Fast-Ion Transport", *Physics Review Letters* 116, 9 (2016).

[11] N. A. Crocker et al. "Three-Wave Interactions between Fast-Ion Modes in the National Spherical Torus Experiment", *Physics Review Letters* 97, 4 (2006).

[12] H. L. Berk et al. "Spontaneous hole-clump pair creation in weakly unstable plasmas", *Physics Letters A* 234, 3 (1997).

[13] M. García-Muñoz et al. "Convective and Diffusive Energetic Particle Losses Induced by Shear Alfvén Waves in the ASDEX Upgrade Tokamak", *Physics Review Letters* 104, 18 (2010).

[14] Y. Kusama et al. "Characteristics of Alfvén eigenmodes, burst modes and chirping modes in the Alfvén frequency range driven by negative ion based neutral beam injection in JT-60U", *Nuclear Fusion* 39, 11Y (1999).

[15] K. Shinohara et al. "Energetic particle physics in JT-60U and JFT-2M", *Plasma Physics and Controlled Fusion* 46, 7 (2004).

[16] N. N. Gorelenkov et al. "The Linear Stability Properties of Medium- to High-n TAEs in ITER", *PPPL Report 4287* (2008).

[17] P. Rodrigues et al. "Systematic linear-stability assessment of Alfvén eigenmodes in the presence of fusion α-particles for ITER-like equilibria", *Nuclear Fusion* 55, 8 (2015).

[18] M. Fitzgerald et al. "Predictive nonlinear studies of TAE-induced alpha-particle transport in Q=10 ITER baseline scenario", *Nuclear Fusion* 56, 11 (2016).

[19] M. Y. Isaev et al. "Alfvén eigenmode evolution computed with the VENUS and KINX codes for the ITER baseline scenario", *Plasma Physics Reports* 43, 2 (2017).

[20] M. Schneller et al. "Nonlinear energetic particle transport in the presence of multiple Alfvénic waves in ITER", *Plasma Physics and Controlled Fusion* 58, 1 (2016).

[21] Y. Todo et al. "Large-Scale Simulation of Energetic Particle Driven Magneto-hydrodynamic Instabilities in ITER Plasmas", *Plasma and Fusion Research* 9, (2014).

[22] R. Waltz et al. "Prediction of the fusion alpha density profile in ITER from local marginal stability to Alfvén eigenmodes", *Nuclear Fusion* 54, 10 (2014).

[23] S. D. Pinches et al. "Energetic ions in ITER plasmas", *Physics of Plasmas* 22, 2 (2015).

[24] A. Mishchenko et al. "Global particle-in-cell simulations of fast-particle effects on shear Alfvén waves", *Physics of Plasmas* 16, 8 (2009).

[25] E. M. Bass et al. "Gyrokinetic simulation of global and local Alfvén eigenmodes driven by energetic particles in a DIII-D discharge", *Physics of Plasmas* 20, 1 (2013).

[26] A. Biancalani et al. "Linear gyrokinetic particle-in-cell simulations of Alfvén instabilities in tokamaks", *Physics of Plasmas* 23, 1 (2016).

[27] M. D. J. Cole et al. "Toroidal Alfvén eigenmodes with nonlinear gyrokinetic and fluid hybrid models", *Physics of Plasmas* 24, 2 (2017).

[28] A. Könies et al. "Benchmark of gyrokinetic, kinetic MHD and gyrofluid codes for the linear calculation of fast particle driven TAE dynamics", *Nuclear Fusion* 58, 12 (2018).

[29] S. Taimourzadeh et al. "Verification and validation of integrated simulation of energetic particles in fusion plasmas", *Nuclear Fusion* 59, 6 (2019).

[30] X. Wang et al. "Nonlinear dynamics of beta-induced Alfvén eigenmode driven by energetic particles", *Physics Review E* 86, (2012).

[31] M. Schneller et al. "Double-resonant fast particle-wave interaction", *Nuclear Fusion* 52, 10 (2012).

[32] S. Briguglio et al. "Particle simulation of bursting Alfvén modes in JT-60U", *Physics of Plasmas* 14, 5 (2007).

[33] A. Bierwage et al. "Role of convective amplification of n=1 energetic particle modes for N-NB ion dynamics in JT-60U", *Nuclear Fusion* 53, 7 (2013).

[34] L. Chen et al. "Physics of Alfvén waves and energetic particles in burning plasmas", *Reviews of Modern Physics* 88, (2016).

[35] P. Lauber. "Super-thermal particles in hot plasmas – Kinetic models, numerical solution strategies, and comparison to tokamak experiments", *Physics Reports* 533, 2 (2013).

[36] B. N. Breizman et al. "Major minority: energetic particles in fusion plasmas", *Plasma Physics and Controlled Fusion* 53, 5 (2011).

[37] W. W. Heidbrink. "Basic physics of Alfvén instabilities driven by energetic particles in toroidally confined plasmas", *Physics of Plasmas* 15, 5 (2008).

[38] A. Fasoli et al. "Chapter 5: Physics of energetic ions", *Nuclear Fusion* 47, 6 (2007).

[39] L. D. Landau. "On the vibration of the electron plasma", *Journal of Physics (U.S.S.R.)* 10, 25 (1946).

[40] T. Hayward. "Multi-mode treatments of the marginally unstable one-dimensional bump-on-tail problem." M.Sc. thesis. Technische Universität München, 2015.

[41] A. A. Vlasov. "On Vibration Properties of Electron Gas", *Journal of Experimental and Theoretical Physics* 3, (1938).

[42] A. J. Brizard et al. "Foundations of nonlinear gyrokinetic theory", *Reviews of Modern Physics* 79, (2007).

[43] S. Possanner. "Gyrokinetics from variational averaging: Existence and error bounds", *Journal of Mathematical Physics* 59, 8 (2018).

[44] E. Zoni et al. "Gyrokinetic theory with polynomial transforms: a model for ions and electrons in maximal ordering", *Journal of Plasma Physics* (2019 (Submitted)).

[45] A. Mishchenko et al. "Mitigation of the cancellation problem in the gyrokinetic particle-in-cell simulations of global electromagnetic modes", *Physics of Plasmas* 24, 8 (2017).

[46] E. Lanti et al. "ORB5: a global electromagnetic gyrokinetic code using the PIC approach in toroidal geometry", *Computer Physics Communications* (2019 Submitted).

[47] J. V. W. Reynders. "Gyrokinetic simulations of finite-β plasmas on parallel architectures." PhD thesis. Princeton University, 1992.

[48] J. C. Cummings. "Gyrokinetic simulations of finite-β and self-generated sheared-flow effects on pressure-gradient-driven instabilities." PhD thesis. Princeton University, 1995.

[49] R. Hatzky et al. "Reduction of the statistical error in electromagnetic gyrokinetic particle-in-cell simulations", *Journal of Plasma Physics* 85, 1 (2019).

[50] R. Hatzky et al. "Electromagnetic gyrokinetic PIC simulation with an adjustable control variates method", *Journal of Computational Physics* 225, 1 (2007).

[51] A. Mishchenko et al. "New variables for gyrokinetic electromagnetic simulations", *Physics of Plasmas* 21, 5 (2014).

[52] A. Mishchenko et al. "Pullback transformation in gyrokinetic electromagnetic simulations", *Physics of Plasmas* 21, 9 (2014).

[53] R. Kleiber et al. "An explicit large time step particle-in-cell scheme for nonlinear gyrokinetic simulations in the electromagnetic regime", *Physics of Plasmas* 23, 3 (2016).

[54] A. Mishchenko et al. "Pullback scheme implementation in ORB5", *Computer Physics Communications* 238, (2019).

[55] R. Fitzpatrick. *Plasma Physics: An Introduction*. CRC Press, Taylor & Francis Group, 2014.

[56] S. D. Pinches. "Nonlinear Interaction of Fast Particles with Alfvén Waves in Tokamaks." PhD thesis. University of Nottingham, 1996.

[57] I. B. Bernstein et al. "An Energy Principle for Hydromagnetic Stability Problems", *Proceedings of the Royal Society of London. Series A* 244, 1236 (1958).

[58] C. Z. Cheng et al. "High-n ideal and resistive shear Alfvén waves in tokamaks", *Annals of Physics* 161, 1 (1985).

[59] R. Betti et al. "Ellipticity induced Alfvén eigenmodes", *Physics of Fluids B: Plasma Physics* 3, 8 (1991).

[60] L. Villard et al. "Geometrical and profile effects on toroidicity and ellipticity induced Alfvén eigenmodes", *Nuclear Fusion* 32, 10 (1992).

[61] R. Betti et al. "Stability of Alfvén gap modes in burning plasmas", *Physics of Fluids B: Plasma Physics* 4, 6 (1992).

[62] N. Miyato et al. "A Modification of the Guiding-Centre Fundamental 1-Form with Strong E×B Flow", *Journal of the Physical Society of Japan* 78, 10 (2009).

[63] C. Z. Cheng. "Kinetic extensions of magnetohydrodynamics for axisymmetric toroidal plasmas", *Physics Reports* 211, 1 (1992).

[64] D. Borba et al. "CASTOR-K: Stability Analysis of Alfvén Eigenmodes in the Presence of Energetic Ions in Tokamaks", *Journal of Computational Physics* 153, 1 (1999).

[65] S. D. Pinches et al. "The HAGIS self-consistent nonlinear wave-particle interaction model", *Computer Physics Communications* 111, (1998).

[66] W. A. Cooper et al. "Full-field drift Hamiltonian particle orbits in 3D geometry", *Plasma Physics and Controlled Fusion* 53, 2 (2011).

[67] D. Spong et al. "Clustered frequency analysis of shear Alfvén modes in stellarators", *Physics of Plasmas* 17, 2 (2010).

[68] A. Könies et al. "From kinetic MHD in stellarators to a fully kinetic description of wave particle interaction", *AIP Conference Proceedings* 1069, 1 (2008).

[69] S. Briguglio et al. "Hybrid magnetohydrodynamic-gyrokinetic simulation of toroidal Alfvén modes", *Physics of Plasmas* 2, 10 (1995).

[70] Y. Todo et al. "Magnetohydrodynamic Vlasov simulation of the toroidal Alfvén eigenmode", *Physics of Plasmas* 2, 7 (1995).

[71] G. Fogaccia et al. "Linear benchmarks between the hybrid codes HYMAGYC and HMGC to study energetic particle driven Alfvénic modes", *Nuclear Fusion* 56, 11 (2016).

[72] G. Y. Fu et al. "Excitation of the toroidicity-induced shear Alfvén eigenmode by fusion alpha particles in an ignited tokamak", *Physics of Fluids B: Plasma Physics* 1, 10 (1989).

[73] R. R. Mett et al. "Kinetic theory of toroidicity-induced Alfvén eigenmodes", *Physics of Fluids B: Plasma Physics* 4, 9 (1992).

[74] J. Candy. "Electron Landau damping of toroidal Alfvén eigenmodes", *Plasma Physics and Controlled Fusion* 38, 6 (1996).

[75] H. L. Berk et al. "Scenarios for the nonlinear evolution of alpha-particle-induced Alfvén wave instability", *Physics Review Letters* 68, (1992).

[76] G. Y. Fu et al. "Stability of the toroidicity-induced Alfvén eigenmode in axisymmetric toroidal equilibria", *Physics of Fluids B: Plasma Physics* 5, 11 (1993).

[77] N. N. Gorelenkov et al. "On the collisional damping of TAE-modes on trapped electrons in tokamaks", *Physica Scripta* 45, 2 (1992).

[78] M. Schneller. "Transport of super-thermal particles and their effect on the stability of global modes in fusion plasmas." PhD thesis. Technische Universität München, 2013.

[79] P. Lauber et al. "LIGKA: A linear gyrokinetic code for the description of background kinetic and fast particle effects on the MHD stability in tokamaks", *Journal of Computational Physics* 226, 1 (2007).

[80] P. Lauber et al. "Analytical finite-Lamor-radius and finite-orbit-width model for the LIGKA code and its application to KGAM and shear Alfvén physics", *Journal of Physics: Conference Series* 1125, (2018).

[81] A. Bierwage et al. "Gyrokinetic analysis of low-n shear Alfvén and ion sound wave spectra in a high-beta tokamak plasma", *Nuclear Fusion* 57, 11 (2017).

[82] F. Zonca et al. "Kinetic theory of low-frequency Alfvén modes in tokamaks", *Plasma Physics and Controlled Fusion* 38, 11 (1996).

[83] P. Lauber et al. "Kinetic Alfvén eigenmodes at ASDEX Upgrade", *Plasma Physics and Controlled Fusion* 51, 12 (2009).

[84] F. Zonca et al. "Existence of discrete modes in an unstable shear Alfvén continuous spectrum", *Plasma Physics and Controlled Fusion* 40, 12 (1998).

[85] J. Candy et al. "Multiplicity of low-shear toroidal Alfvén eigenmodes", *Physics Letters A* 215, 5 (1996).

[86] S. Jolliet et al. "A global collisionless PIC code in magnetic coordinates", *Computer Physics Communications* 177, 5 (2007).

[87] A. Bottino et al. "Global simulations of tokamak microturbulence: finite-β effects and collisions", *Plasma Physics and Controlled Fusion* 53, 12 (2011).

[88] N. Tronko et al. "Second order gyrokinetic theory for particle-in-cell codes", *Physics of Plasmas* 23, 8 (2016).

[89] A. Bottino. "Monte Carlo particle-in-cell methods for the simulation of the Vlasov-Maxwell gyrokinetic equations", *Journal of Plasma Physics* (2019 (in prep.)).

[90] R. C. Davidson. *Methods in Nonlinear Plasma Theory*. Vol. 37. Pure and Applied Physics. Academic Press, 1972.

[91] A. Biancalani et al. "Interaction of Alfvénic modes and turbluence, investigated in a self-consistent gyrokinetic framework," (Milan, July 12, 2019). 2019.

[92] N. Ohana et al. "Gyrokinetic Simulations on Many- and Multi-core Architectures with the Global Electromagnetic Particle-In-Cell Code ORB5", *Computer Physics Communications* (2019 Submitted).

[93] I. Novikau et al. "Implementation of energy transfer technique in ORB5 to study collisionless wave-particle interactions in phase-space", *Computer Physics Communications* (2019 Submitted).

[94] A. Figueiredo et al. "Comprehensive evaluation of the linear stability of Alfvén eigenmodes driven by alpha particles in an ITER baseline scenario", *Nuclear Fusion* 56, 7 (2016).

[95] P. Rodrigues et al. "Sensitivity of alpha-particle-driven Alfvén eigenmodes to q-profile variation in ITER scenarios", *Nuclear Fusion* 56, 11 (2016).

[96] C. Angioni et al. "Gyrokinetic simulations of impurity, He ash and α particle transport and consequences on ITER transport modelling", *Nuclear Fusion* 49, 5 (2009).

[97] Y. Wu et al. "Nonlinear evolution of the alpha-particle-driven toroidicity-induced Alfvén eigenmode", *Physics of Plasmas* 2, 12 (1995).

[98] N. N. Gorelenkov et al. "Resonance broadened quasi-linear (RBQ) model for fast ion distribution relaxation due to Alfvénic eigenmodes", *Nuclear Fusion* 58, 12 (2018).

[99] S. D. Pinches et al. "Spectroscopic determination of the internal amplitude of frequency sweeping TAE", *Plasma Physics and Controlled Fusion* 46, 7 (2004).

[100] G. Y. Fu. "Existence of core localized toroidicity-induced Alfvén eigenmode", *Physics of Plasmas* 2, 4 (1995).

[101] H. L. Berk et al. "More on core-localized toroidal Alfvén eigenmodes", *Physics of Plasmas* 2, 9 (1995).

[102] S. Sharapov et al. “Stability of alpha particle driven Alfvén eigenmodes in high performance JET DT plasmas”, *Nuclear Fusion* 39, 3 (1999).

[103] B. J. Tobias et al. “Fast Ion Induced Shearing of 2D Alfvén Eigenmodes Measured by Electron Cyclotron Emission Imaging”, *Physics Review Letters* 106, 7 (2011).

[104] Z. X. Lu et al. “Kinetic effects of thermal ions and energetic particles on discrete kinetic BAE mode generation and symmetry breaking”, *Nuclear Fusion* 58, 8 (2018).

Acknowledgments

I would like to thank some of the many people who have helped me throughout my work on this thesis.

First of all, I am indebted to Dr. Philipp Lauber, my supervisor, for his insight and assistance. His expertise on the subject at hand is unquestionable and his supervision has fostered my own devlopment, through discussions about all aspects of the thesis.

Secondly, I would like to thank Prof. Sibylle Günter, my academic supervisor and division head, for making it possible to pursue the topic. I am especially thankful that, despite her busy schedule, she was able to provide feedback and suggestions both for this document and in discussions throughout my research.

Next, I would like to explicitly acknowledge those who at times acted as unofficial supervisors, Dr. Alberto Bottino, with whom I had many discussions about the simulations with, and model and details of, the ORB5 code. I also thank Dr. Xin Wang, with whom I discussed physics and results constantly, and who never failed to answer my stupid questions. Similarly, Dr. Zhixin Lu, with whom have I discussed many physical and numerical issues.

I thank Dr. Simon Pinches, now at ITER, whose code HAGIS formed a significant part of my work, also for his jovial nature at conferences, as well as for hosting an inspiring short visit to ITER.

On the side of the ORB5 code, in addition to Alberto, I give thanks to the team of people who have been working on the code, in particular to Noé Ohana and Dr. Emmanuel Lanti of EPFL for their work on the modernisation of the code and development of a test framework. Finally, such global electromagnetic simulations as performed in this thesis would not have been possible without the work by Alberto, Dr. Roman Hatzky, and our colleagues in Greifswald, especially Dr. Alexey Mishchenko, for their progress on mitigating the cancellation problem.

Within the wider institute, I would like to thank a few colleagues, Prof. Emanuele Poli, briefly my acting division head when I started, for showing interest in the results we obtained during this work, also presenting them at conferences in a positive light. I am grateful to Dr. Dave Coster, the *de facto* local expert in all matters numerical, to Dr. Bruce Scott, whose great wealth of knowledge I consulted on several occasions, and to Dr. Geri Papp, to whom I would turn when I had issues with signal processing. I would like to thank my office mates over my time at IPP, in particular the two people

with whom I shared an office the longest, Nils Leuthold and Dr. Branka Vanovac, with whom many jokes and many cups of coffee were shared.

Finally, special thanks to my family for their support, in particular to my wonderful wife, Stephanie.