Thomas W. Hayward

Multi-mode treatments of the marginally unstable one-dimensional bump-on-tail problem

Max-Planck-Institut
für Plasmaphysik

Bereich Tokamakphysik

Technical
University
of Munich

Fakultät für Physik

MSc Thesis

Multi-mode treatments of the marginally unstable one-dimensional bump-on-tail problem

Thomas W. Hayward
September 2015

Contents

Abstract

Fusion alpha particles or neutral beam heated plasma ions can provide resonant particle drive to Alfvénic instabilities in magnetic confinement fusion devices. Close to marginal stability, these instabilities can be reduced to the simpler one-dimensional bump-on-tail system via a mapping along the resonance, whilst retaining a significant amount of the physical behaviour.

Collisionless Vlasov-Poisson simulations of the electron bump-on-tail system have been performed with both Eulerian and particle based tools. First a series of benchmarks have been performed with single mode simulations, and with strongly driven multi-mode simulations. Then marginally unstable multi-mode simulations were performed and a transport threshold has been observed above which phase space redistribution is enhanced as a result of the overlap of phase space structures. An insensitivity of this effect to the relative initial amplitudes has been demonstrated in the case of a two-mode system. We show that the ability to track particle trajectories using the particle method can be a powerful tool for mapping the phase space position of a mode, which allows us to make the claim that the threshold of enhanced transport corresponds to the intersection in phase space of different modes.

It is also observed that Eulerian and particle methods have very different behaviour as the resolution is decreased below what is sufficient. Eulerian methods suffer spurious diffusion when the grid cells are no longer sufficiently fine to resolve the scales of the mode. In contrast, we see that particle methods tend to under predict the transport as the number of markers is decreased. We postulate that if particle methods are missing the critical trajectories, through which the modes overlap, then the mechanism of the interaction is missed, and the simulation lacks the transport.

Zusammenfassung

Bei der Fusion mittels magnetischen Einschlusses befinden sich im Reaktor sowohl alpha-Teilchen der Fusion als auch schnelle Ionen der Neutralteilchenheizung. Diese Teilchen können resonant mit Alfvénischen Instabilitäten wechselwirken. Nahe der marginalen Stabilität kann die Beschreibung dieser Instabilitäten auf ein einfacheres, eindimensionales "bump-on-tail" Modell reduziert werden. Dies geschieht durch das Abbilden des Systems auf die Resonanzregion, wobei viele physikalische Eigenschaften erhalten bleiben.

Stoßfreie Vlasov-Poisson Simulationen des Elektronen bump-on-tail Modells wurden sowohl mit Euler- als auch mit Particle-in-Cell basierenden Methoden ausgeführt. Zunächst wurde eine Reihe von Benchmark-Tests mit Einzelmoden-Simulationen durchgeführt, sowie auch mit stark getriebenen Multimoden-Simulationen. Danach wurde in marginal instabilen Multimoden-Simulationen ein Schwellwert für den Transport beobachtet. Oberhalb dieses Schwellwerts findet als Folge überlappender Phasenraumstrukturen eine verstärkte Umverteilung im Phasenraum statt. Die Unabhängigkeit dieses Effekts von der Anfangsamplitude konnte für ein Zwei-Moden-System demonstriert werden. Es wird gezeigt, dass die Möglichkeit der Particle-in-Cell Methode Teilchentrajektorien zu verfolgen ein mächtiges Werkzeug sein kann um die Phasenraumposition der Mode abzubilden. Damit konnte festgestellt werden, dass der verstärkte Transport mit dem Überschneiden von verschiedenen Moden im Phasenraum einhergeht.

Außerdem wurde beobachtet, dass sich das Verhalten von Euler- und Particle-in-Cell-Methoden stark voneinander unterscheidet, wenn die Auflösung die nötigen Werte unterschreitet. Euler-Methoden zeigen unphysikalische Diffusion, wenn die Zellen des Simulationsgitters nicht mehr klein genug sind, um die Strukturen der Mode aufzulösen. Im Gegensatz dazu führen Particle-in-Cell-Methoden zu einer Unterschätzung des Transportes wenn die Zahl der Marker zu stark reduziert wird. Man kann die Behauptung aufstellen, dass in Particle-in-Cell-Methoden das Fehlen der kritischen Trajektorien, welche den Modenüberlapp ermöglichen dazu führt, dass der Wechselwirkungsmechanismus nicht beschrieben werden kann, woraufhin die Simulation den dazugehörigen Transport verfehlt.

Chapter 1

Introduction and Motivation

A burning deuterium-tritium (D-T) fueled magnetic fusion reactor will produce energy via the fusion of a deuterium and a tritium nucleus.

$$D + T \rightarrow {}^4He + n$$
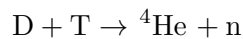
The maximum reaction rate occurs when the reactants have energies of $70\,\mathrm{keV}$ [1], however, because of the hot tail of a Maxwellian distribution, a future fusion reactor is likely to have a plasma temperature of 10–$20\,\mathrm{keV}$ [2].

The neutron released from the reaction has $14\,\mathrm{MeV}$ energy and leaves the plasma without interaction, where its energy must then be captured and converted into electricity. The helium nucleus (also known as an alpha particle), with energy $3.6\,\mathrm{MeV}$, as a charged particle, is bound by the magnetic field of the reactor, and therefore remains in the plasma, where it can be used as a means of heating the plasma. As the population of alpha particles have energies much in excess of those of the thermal background plasma species, there exists the possibility that this species may also interact with electromagnetic modes in the plasma, driving kinetic energetic particle instabilities [3].

Even in a non-burning plasma, either a present day externally heated scientific experiment or the initial heating phase of a future reactor, which are typically heated via beams of high energy neutral particles which become ionized upon entry into the plasma, and via resonant heating of ions and electrons using radio waves, a fast ion species can already trigger the onset of energetic particle driven instabilities [4].

One of the relevant types of magnetic plasma waves are the so-called Alfvén waves, discovered by H. Alfvén [5]. A detailed introduction to the physical nature of Alfvén waves is presented by L. Chen [6]. As these instabilities have the potential [7] to lead to an increased number of particle losses from the plasma to the plasma wall, and to transport particles and energy away from the hot plasma core, it is important to study the processes governing the dynamics of the instabilities and the coupling between the fast particles present and the electromagnetic modes.

The interaction between an energetic particle species and a single electromagnetic mode has been studied in the past, and the drive of the mode is dominated, at least in the linear phase, by those particles with velocities equal to the phase velocity of the wave. As this interaction is based on a velocity resonance, we expect that the changes to the plasma take place in a limited part of phase space close to the resonance, and will only affect the remainder of phase space if the perturbed distribution becomes large. We shall therefore use a kinetic model for the system, as fluid models are unable to account for the physics of resonant interactions.

The tokamak [2], a name which comes from the Russian for "Toroidal chamber with magnetic coils," is a fusion reactor design concept. The shape of the plasma is toroidal,

and the plasma is confined with a helical magnetic field — the toroidal component of which is generated externally by magnetic coils, and the poloidal component is generated via induction by driving a current through the plasma itself, which is made possible by treating the plasma as the secondary coil of a transformer. Given the toroidal periodicity of the device, we can consider a toroidal mode number, n, to be the number of toroidal rotations which are undergone before mapping to the same point — and similarly there exists a poloidal mode number, m, which is the number of poloidal rotations which are undergone. Experimental observations in many present-day tokamaks [8, 9] reveal Alfvénic behaviour with strong nonlinearities.

Whether individual modes of the system are stable or unstable depends on the plasma equilibrium. Whereas in ASDEX Upgrade (AUG), in discharges with Alfvénic plasmas, unstable modes are seen with toroidal mode number $\lesssim 5$ [10], in ITER, typical mode numbers predicted to be unstable are of the order of $15 \lesssim n \lesssim 35$ with significant overlapping of mode structure [11, 12]. This suggests that for modelling of ITER or reactor-type devices, we can no longer assume that we can treat each Alfvénic mode independently.

As discussed in [13], when below a threshold, the transport caused by a mode scales with the radial amplitude of the magnetic field perturbation $\sim \delta B_r/B_0$, but above the threshold, scales quadratically with the radial amplitude $\sim \left(\delta B_r/B_0\right)^2$. For a single mode, this threshold is of the order of $\delta B_r/B_0 = 10^{-3}$. However, when multiple modes are present, this threshold is reduced to $\delta B_r/B_0 = 10^{-4}$.

When dealing with the nonlinear behaviour exhibited in the system, nonlinear numerical simulation plays an important part in the understanding of complex plasma systems. Simulations are used as a tool for the modelling and understanding of experimental data, especially in showing which of competing mechanisms is responsible for observed behaviour. It can also be used as a synthetic diagnostic, as a means for calculating certain experimental plasma parameters which perhaps cannot be measured directly. Finally, simulation plays an important part in making predictions, both for the behaviour of possible scenarios of existing machines and for the performance of, and therefore in the design of, future machines.

Berk and Breizman [14] show that close to marginal stability the Alfvénic system can be equated via a change of variables and a mapping along the resonance line to the one-dimensional bump-on-tail problem, which allows a reduction of the dimensionality of the problem from six- to two-dimensional phase space. Within the framework of the one-dimensional bump-on-tail problem, linear theory is capable of predicting the linear growth rates of the modes, and quasi-linear theory approximates the mode saturation amplitude for a single mode. Berk et al. extend [15, 16] this quasi-linear theory to predict an increase in the saturation amplitudes of a system with multiple modes if the modes overlap, by assuming that phase space is flattened across the union of the regions flattened in the case of single isolated modes.

Previous works [17, 18, 19] present numerical results of Berk and Breizman's augmented bump-on-tail model, in which a field damping term γ_d has been added to the Vlasov-Ampère system of equations, along with a collision operator, and have focussed on the characterization of nonlinear behaviour of single modes.

The interaction of multiple Alfvénic resonant modes with fast particles has been simulated using a hybrid model with Hagis and Ligka [20]. Some of the findings, namely the increased nonlinear saturation level caused by the flattening of the entire overlapping phase space region, had been predicted by quasi-linear theory, however other aspects, namely that the linearly subdominant mode can become nonlinearly dominant, cannot be predicted be quasi-linear theory, which assumes phase space diffusion in the

2

case of overlapping modes, and therefore omits any effects which rely on the existence and preservation of coherent fine-scale structures in phase space.

When wanting to perform numerical simulations of fast particle instabilities, an interesting question is how the numerical requirements will scale with the physical parameters, and how this differs between particle and grid based methods. It is surmised that grid based numerical simulation requires a grid capable of resolving the structures in phase space, particularly in the velocity dimensions. On the other hand, it is somewhat more unclear how the requirements of particle based methods scale in terms of the physical scales. The question of how these requirements scale in the case of interacting multiple modes, in which phase space may become partially diffusive, is not well understood and this is one of the problems that we seek to address. It is also important to understand how the results from simulations fail as the resolution is reduced below the required level, and to understand why this is so.

In this work, we investigate both single- and multi-mode collisionless bump-on-tail systems, looking at behaviour both close to ($\gamma/\omega < 1\%$) and far from ($\gamma/\omega \approx 20\%$) marginal stability. We perform nonlinear Vlasov-Poisson simulations using both Eulerian (grid based) and PIC (particle based) methods. We pay particular attention to the growth rates and saturation amplitudes, and the different numerical challenges of obtaining these data with two different numerical tools.

Chapter 2

Theoretical background

If we consider a one-dimensional system of plasma particles, then we can think of this system as containing a two-dimensional phase space, where each particle (index i) is identified by a position (x_i), and a velocity (v_i).

We can then consider $f(x, v, t)$ to be the density of particles with phase space coordinates (x, v) at time t, which we shall refer to as the distribution function:

$$f(x, v, t) = \sum_i m_i \delta(x - x'_i(t)) \delta(v - v'_i(t)) \tag{2.1}$$

If we take the total derivative of distribution function to be zero,

$$\frac{\mathrm{d}f}{\mathrm{d}t} = \frac{\partial f}{\partial t} + \frac{\mathrm{d}x}{\mathrm{d}t} \cdot \frac{\partial f}{\partial x} + \frac{\mathrm{d}v}{\mathrm{d}t} \cdot \frac{\partial f}{\partial v} = 0$$

and using that

$$\frac{\mathrm{d}x}{\mathrm{d}t} \equiv v \qquad\qquad \frac{\mathrm{d}v}{\mathrm{d}t} \equiv a = \frac{q}{m}(E + v \times B)$$

we arrive at the Vlasov equation [21], which is the collisionless Boltzmann equation,

$$\frac{\partial f}{\partial t} + v \cdot \frac{\partial f}{\partial x} + \frac{q}{m}(E + v \times B) \cdot \frac{\partial f}{\partial v} = 0 \tag{2.2}$$

which, when combined with Maxwell's equations,

$$\nabla \times E = -\frac{\partial B}{\partial t} \qquad\qquad \nabla \times B = \mu_0 (J + \epsilon_0 \frac{\partial E}{\partial t})$$

$$\nabla \cdot E = \frac{\rho}{\epsilon_0} \qquad\qquad \nabla \cdot B = 0$$

defines the Vlasov-Maxwell system.

In the special case of a non-relativistic system, and where the magnetic field varies only slowly with time, we can decouple the electric and magnetic fields, and neglect the $v \times B$ term in the Lorentz force. Then, we arrive at the electrostatic Vlasov equation,

$$\frac{\partial f}{\partial t} + v \cdot \frac{\partial f}{\partial x} + \frac{q}{m} E \cdot \frac{\partial f}{\partial v} = 0 \tag{2.3}$$

which we can couple with

$$\nabla \times E = 0 \tag{2.4}$$

and
$$\nabla \cdot E = \frac{\rho}{\epsilon_0} \tag{2.5}$$

to give us the Vlasov-Poisson system, noting that equation 2.4 allows us to write the electric field E in terms of a scalar potential φ,

$$E = -\nabla\varphi$$

$$\nabla^2\varphi = -\frac{\rho}{\epsilon_0} \tag{2.6}$$

$$\rho(x,t) = q_{\mathrm{i}} \int f_{\mathrm{i}}(x,v,t)\,\mathrm{d}v + q_{\mathrm{e}} \int f_{\mathrm{e}}(x,v,t)\,\mathrm{d}v$$

As we want to focus only on the fast dynamics of the system, we can assume that the ions are motionless and represent a constant uniform neutralizing distribution with overall quasi-neutrality,

$$f_{\mathrm{i}} = \mathrm{const} \qquad\qquad \int q_{\mathrm{i}} n_{\mathrm{i}}\,\mathrm{d}x + \int q_{\mathrm{e}} n_{\mathrm{e}}\,\mathrm{d}x = 0$$

$$q_{\mathrm{i}} = e \qquad\qquad q_{\mathrm{e}} = -e$$

$$\int n_{\mathrm{i}}(x,t)\,\mathrm{d}x = \int n_{\mathrm{e}}(x,t)\,\mathrm{d}x = n$$

therefore,

$$\rho(x,t) = en\left(1 - \frac{n_{\mathrm{e}}(x,t)}{n}\right)$$

We choose to expand the electron distribution function into a constant and a spatially and temporally varying part,

$$f(x,v,t) = f_0(v) + f_1(x,v,t)$$

noting that

$$\frac{\partial f_0}{\partial t} = 0 \qquad\qquad\qquad \frac{\partial f_0}{\partial x} = 0$$

Now, if we substitute these expansions into the Vlasov equation,

$$\frac{\partial f_1}{\partial t} + v \cdot \frac{\partial f_1}{\partial x} - \frac{e}{m_{\mathrm{e}}}\left(E_1 \cdot \frac{\partial f_0}{\partial v} + E_1 \cdot \frac{\partial f_1}{\partial v}\right) = 0 \tag{2.7}$$

We can also substitute this into the charge density equation, using

$$n_{\mathrm{e}} = \int f\,\mathrm{d}v = \int f_0\,\mathrm{d}v + \int f_1\,\mathrm{d}v = n + \int f_1\,\mathrm{d}v$$

$$\rho(x,t) = e\int f_1(x,v,t)\,\mathrm{d}v$$

which allows us to write

$$\nabla^2\varphi = \frac{e}{\epsilon_0}\int f_1(x,v,t)\,\mathrm{d}v \tag{2.8}$$

This gives us the final nonlinear electrostatic Vlasov-Poisson system, under the assumptions of fixed background ions, fully determined by equations 2.7 and 2.8.

Linear reponse

If we look carefully at the Vlasov-Poisson equations, we see that as $E_1 = -\nabla \varphi$ depends on f_1 in the Poisson equation, and these dependent quantities (E_1 and $\frac{\partial f_1}{\partial v}$) both appear together in the final term of the left-hand side (LHS) of the Vlasov equation. This is the nonlinearity in the Vlasov-Poisson system.

If we want to find out how the plasma would react to a small perturbation, we can first perform a linearization of the Vlasov-Poisson equations. In order to do this, we must first discuss an ordering of the terms. If we restrict ourselves to small perturbations, then we can assume that the spatial and temporal part of the distribution function is small in comparison to the constant part, $f_0(v) \gg f_1(x, v, t)$, and we introduce a small parameter, $f_1/f_0 \sim \delta \ll 1$. We also note that the electric field due to f_0 is zero, and the electric field is only due to f_1 and is therefore also small. This is the reason for our inclusion of the subscript "1" on the electric field earlier. Therefore, we argue that the nonlinear term, $E_1(x, t)\frac{\partial f_1}{\partial v}$ is small compared with $E_1(x, t)\frac{\partial f_0}{\partial v}$ and the other terms in the Vlasov equation, which themselves are small, and we neglect this $\sim \delta^2$ term.

This allows us to write a linearized Vlasov equation,

$$\frac{\partial f_1}{\partial t} + v \cdot \frac{\partial f_1}{\partial x} - \frac{e}{m_{\mathrm{e}}} E_1 \frac{\partial f_0}{\partial v} = 0 \tag{2.9}$$

From this linearized equation, we can derive a dispersion relation linking the wave number, k, and the complex frequency, ω, whose imaginary part, $\gamma = \mathbb{Im}\{\omega\}$ governs the growth rate or decay rate of the perturbation. This procedure was first done by Landau [22], who improved on and corrected the work of Vlasov [21]. We shall not present the full details of the derivation of the dispersion relation, which can be read in Landau's paper, or in a number of other guides (e.g. Stix [23], Brambilla [24], Fitzpatrick [25]), but rather we summarize the key results following the works of Landau and Fitzpatrick, and attempt to provide the arguments justifying some of the important steps.

Although we are interested in the results for the one-dimensional system in which $\mathbf{v} = v$, the derivation was originally performed for the more general system in which $\mathbf{v} = \{v_x, v_y, v_z\}$. Therefore, Landau introduced the variable u for the parallel velocity $u = k \cdot \mathbf{v}$. Although this is not necessary for us, we shall nevertheless follow this notation.

Landau [22] tells us that we may not perform a Fourier transform in time, but instead we must treat this as an initial value problem. We therefore perform a Laplace transform in time, and Fourier transform in space, using the Laplace transform with parameter p,

$$\bar{f}_1(u, p) = \int_0^\infty f_1(u, t)e^{-pt}\, \mathrm{d}t \tag{2.10}$$

which has an inverse transform,

$$f_1(u, t) = \frac{1}{2\pi i} \int_C \bar{f}_1(u, p)e^{pt}\, \mathrm{d}p \tag{2.11}$$

and noting that the Laplace transform of $\frac{\partial f_1}{\partial t}$ is $p\bar{f}_1 - f_1(u, t = 0)$,

$$ik\bar{E}_1 = \frac{e}{\epsilon_0} \int f_1(u)\, \mathrm{d}u$$

$$p\bar{f}_1 + iku\bar{f}_1 = \frac{e}{m_{\mathrm{e}}} \bar{E}_1 \frac{\partial f_0}{\partial u} + f_1(u, t = 0)$$

7

$$\bar{E}_1 = -\frac{e/\epsilon_0}{ik\epsilon(k,p)} \int_{-\infty}^{\infty} \frac{f_1(u,t=0)}{p+iku} \, du$$

with

$$\epsilon(k,p) = 1 + \frac{e^2}{\epsilon_0 m_e k} \int_{-\infty}^{\infty} \frac{\partial f_0/\partial u}{ip - ku} \, du$$

allowing us to arrive at the Laplace transform of the perturbed distribution function,

$$\bar{f}_1(u,p) = \frac{e}{m_e}\bar{E}_1 \frac{\partial f_0/\partial u}{p+iku} + \frac{f_1(u,t=0)}{p+iku}$$

We can then invert the Laplace transform using equation 2.11 after a careful consideration of the complex path C. We choose to displace the path, which was originally running vertically to the right of the imaginary axis (as described by the inverse Laplace transform), into the left half plane, yet it must be deformed around the poles rather than crossing them. We find the contribution from the vertical part of the contour is exponentially smaller than the residue from the right-most pole in the plane, which is the only important contribution to the integral, and therefore the long term behaviour of the perturbation and the electric field is dominated by this pole.

We can therefore revert the dispersion relation in terms of ω rather than p, and by using the same reasoning as for p, the poles of u must be assumed to be in positive-imaginary half plane, and then analytically continued into the negative half plane, by deforming the contour around the poles. This gives us the dispersion relation,

$$\epsilon(k,\omega) = 1 + \frac{e^2}{\epsilon_0 m_e k} \int_{-\infty}^{\infty} \frac{\partial f_0/\partial u}{\omega - ku} \, du = 0 \tag{2.12}$$

and we note the pole in the denominator when $u = \omega/k$. If we assume that u is mostly real at the pole, or $\mathbb{Re}\{\omega\} \gg \gamma$, then

$$\int_{-\infty}^{\infty} \frac{\partial f_0/\partial u}{\omega - ku} \, du = P \int_{-\infty}^{\infty} \frac{\partial f_0/\partial u}{\omega - ku} \, du - \frac{i\pi}{k} \left.\frac{\partial f_0}{\partial u}\right|_{u=\omega/k} \tag{2.13}$$

where $P \int$ denotes the principal value of the integral,

$$P \int_{-\infty}^{\infty} \frac{(\cdots)}{x-a} \, dx = \lim_{\epsilon \to 0^+} \int_{-\infty}^{a-\epsilon} \frac{(\cdots)}{x-a} \, dx + \int_{a+\epsilon}^{\infty} \frac{(\cdots)}{x-a} \, dx$$

If we assume that our background distribution function is that of a Maxwellian distribution, given by

$$f_0(v) = n\sqrt{\frac{m_e}{2\pi T_e}} \exp\left(\frac{-v^2}{2T_e m_e}\right) \tag{2.14}$$

then, in the case where k is small, we can perform an expansion of the denominator of equation 2.13 using the scaling $\omega \gg ku$,

$$\frac{1}{\omega - ku} \approx \frac{1}{\omega}\left(1 + \frac{ku}{\omega} + \frac{k^2 u^2}{\omega^2} + \frac{k^3 u^3}{\omega^3} + \cdots\right)$$

integrating term by term leads us to the approximation for the real frequency,

$$\omega^2 \approx \omega_p^2\left(1 + 3k^2\lambda_D^2\right) \tag{2.15}$$

8

with $\omega_{\mathrm{p}} = \sqrt{ne^2/\epsilon_0 m_{\mathrm{e}}}$ and $\lambda_{\mathrm{D}} = \sqrt{T_{\mathrm{e}}/m_{\mathrm{e}}\omega_{\mathrm{p}}^2}$. If we assume that γ (the imaginary part of ω) is small, such that we can expand $\omega = \omega_0 + \gamma$, taking ω_0 as the solution to equation 2.15, then we can write

$$\gamma \approx \frac{i\pi}{2} \frac{e^2 \omega_{\mathrm{p}}}{\epsilon_0 m_{\mathrm{e}} k^2} \left. \frac{\partial f_0}{\partial u} \right|_{u=\omega/k} \tag{2.16}$$

and if we substitute f_0 from equation 2.14, then we arrive at the expression for the Landau damping decay rate,

$$\gamma \approx -\frac{i}{2} \sqrt{\frac{\pi}{2}} \frac{\omega_{\mathrm{p}}}{(k\lambda_{\mathrm{D}})^3} \exp\left(-\frac{1}{2(k\lambda_{\mathrm{D}})^2}\right) \tag{2.17}$$

Growing modes

As we have seen in equation 2.16, the growth rate (or damping rate) is approximately proportional to the gradient of the distribution function (f_0) with respect to the velocity (v or u), as evaluated at the phase velocity of the perturbation, $\gamma \propto \left. \partial f_0/\partial v \right|_{v=\omega/k}$. In equation 2.17, we then substituted in the distribution function f_0 from equation 2.14, and we see that this gave us damping.

However, if our distribution function is not simply a Maxwellian distribution, but rather the sum of Maxwellian distributions, as in equation 2.18,

$$f_0(v) = n \left(\sum_{\mathrm{s}} n_{\mathrm{s}} \sqrt{\frac{m_{\mathrm{e}}}{2\pi T_{\mathrm{s}}}} \exp\left(-\frac{(v-v_{\mathrm{s}})^2}{2T_{\mathrm{s}} m_{\mathrm{e}}}\right) \right) \tag{2.18}$$

where $\sum n_{\mathrm{s}} = 1$, then if there exists one $v_{\mathrm{s}} \neq 0$, it is no longer guaranteed that $\left. \partial f_0/\partial v \right|_{v=\omega/k}$ is negative.

We interest ourselves in the bump-on-tail system, where we have two species, the main plasma or bulk plasma (s = p) with no velocity offset $v_{\mathrm{p}} = 0$, nominal temperature $T_{\mathrm{p}} = T$, which comprises the majority of the total density $n_{\mathrm{p}} = 1 - n_{\mathrm{b}} \gg n_{\mathrm{b}}$. In addition, we have a minority beam plasma (s = b), which has a non-zero velocity offset $v_{\mathrm{b}} > 0$, a temperature T_{b}, and a density n_{b}.

In our typical cases, where $\frac{1}{4}T \lesssim T_{\mathrm{b}} \lesssim T$, $v_{\mathrm{b}} \approx 5v_{\mathrm{th}}$, and $n_{\mathrm{b}} \lesssim 0.1$, we indeed find that for a range of values of k, $\gamma > 0$.

Saturation

All of the results outlined in this chapter have relied on the linearization of the Vlasov equation, which we performed by neglecting the term $\frac{e}{m_{\mathrm{e}}} E_1 \frac{\partial f_1}{\partial v}$ in the Vlasov equation, under the justification that $f_0 \gg f_1$. In the case of Landau damping, or any other decaying modes, it is sufficient to initialize the mode with a sufficiently small amplitude in order to keep this justification valid. However, if the mode is a growing mode, then the linear result predicts continual exponential growth of the mode amplitude, and therefore of f_1. Therefore, in the case of growing modes, the assumption of $f_0 \gg f_1$ will always become false at sufficiently large times.

What we find is that as the amplitude of the perturbation becomes large, particles become trapped in the potential well, and the perturbed distribution function flattens the positive gradient in the initial distribution function, thus reducing and eventually stopping the mode growth. We can also think of the gradient inversion in distribution

function of the bump-on-tail system as a source of free energy. Therefore, if a mode grows, and therefore energy is transferred to the electric field, then energy must come from the kinetic energy of the particles in the distribution function.

Particle Trapping

We can consider the potential under the assumption of a single harmonic, $\varphi(x,t) = \varphi_{\mathrm{amp}}(t) \cdot \varphi_{\mathrm{phase}}(x,t)$, as a time-dependent amplitude part

$$\varphi_{\mathrm{amp}}(t) = \varphi_{\mathrm{amp}}(t=0) \cdot \exp\left(\gamma t\right)$$

and a phase part,

$$\varphi_{\mathrm{phase}}(x,t) = \exp\left(i\left(kx - \mathrm{Re}\{\omega\}t\right)\right)$$

which depend on the imaginary and the real parts of the frequency respectively.

The potential is moving with a phase velocity $v_{\mathrm{p}} = \frac{\mathrm{Re}\{\omega\}}{k}$, and we can write the phase of the potential,

$$\varphi_{\mathrm{phase}}(x,t) = \exp\left(ik\left(x - v_{\mathrm{p}}t\right)\right)$$

We can therefore move into the frame of the wave if we replace the coordinate v with $v - v_{\mathrm{p}}$, such that the phase of the potential is now fixed.

At the base of the potential well, we see that the kinetic energy of the particles is below the potential of the well,

$$W = \frac{m_{\mathrm{e}}(v - v_{\mathrm{p}})^2}{2} - e\varphi_{\mathrm{amp}}(t)\cos\left(kx\right) < 0$$

and the particles are therefore trapped. At the threshold of a particle being trapped in the period of the wave, we instead replace the inequality above with an equality,

$$W = \frac{m_{\mathrm{e}}(v - v_{\mathrm{p}})^2}{2} - e\varphi_{\mathrm{amp}}(t)\cos\left(kx\right) = 0$$

which defines an eliptical orbit in phase space with real space semi axis $\Delta x = \frac{\pi}{2k}$ and velocity space semi axis $\Delta v = \sqrt{\frac{2e}{m_{\mathrm{e}}}\varphi_{\mathrm{amp}}}$. All particles with phase space positions inside this separatrix are said to be trapped as their orbits remain confined within the separatrix, and they therefore only encounter a single period of the wave. We assume that the asymptotic behaviour leads to the distribution function becoming constant along a phase space trajectory. We also assume that as the growth of the mode is slow and continuous, that there is no fine scale structure within the island. These assumptions lead to phase space becoming flattened across the whole of the island, and we assume that $f(x, v, t \to \infty)_{\mathrm{island}} = f_0(v_{\mathrm{p}})$ [26].

In the quasi-linear model [16] the assumptions are yet stronger, assuming that the distribution function in all of phase space between $v_- = v_{\mathrm{p}} - \frac{\Delta v}{2}$ and $v_+ = v_{\mathrm{p}} + \frac{\Delta v}{2}$ is flattened. By equating the change in particle momentum with a change in wave momentum, which depends on both the width of the flattened region and the gradient of the distribution function across that region, an expression can be derived showing a quadratic dependence of the mode saturation amplitude and the mode growth rate.

$$\omega_{\mathrm{b}} \equiv \left|\frac{2e\varphi k^2}{m_{\mathrm{e}}}\right|^2 \approx \frac{8}{3\pi}\gamma$$

Multiple modes

As discussed in [16], in which multi-mode results are presented for the two-steam instability using a quasi-linear model, the assumption of the distribution function becoming flattened across the region within the island's separatrix has interesting implications if there are multiple modes whose velocity space domains overlap. It provides a mechanism for an enlarged region of phase space, namely the union of the domains of a chain of overlapping islands, to undergo collective phase space flattening, which implies a much increased level of the saturation amplitude. We shall investigate this effect in §4.4.

Chapter 3

Numerical tools

3.1 Introduction

In general, the Vlasov-Poisson set of equations, when combined with a set of initial conditions and boundary conditions, entirely dictate the behaviour of the system in question. However, as these equations are nonlinear, there is seldom an analytic expression for the long term evolution or for the final state of the system. Therefore it is almost always necessary to turn to nonlinear numerical simulation, which iteratively evolves the system using the equations and boundary conditions laid out.

In general, there are two main approaches for the design of a numerical code to solve our set of equations. Both govern methods for discretizing the equations in such as way as they can be computed numerically. We denote the two methods as Vlasov (§3.2) and Particle-in-Cell (PIC) (§3.3) methods.

An important difference between Vlasov and PIC methods is that, whilst some numerical schemes for Vlasov methods might show better convergence in one-dimension than PIC methods, any full grid Vlasov method is subject to the so-called "Curse of dimensionality", whereby schemes experience a decrease in their convergence properties proportional to the number of dimensions. This is related to the fact that, as the number of points in each dimension of a grid is doubled, the total number of points in the grid increases with 2^d. However, as the PIC method is a type of Monte Carlo method, the error scales independently of the number of dimensions. For this reason, even though in our two-dimensional phase space system, Vlasov methods offer very good properties and performance, as one turns to higher dimensional systems (for example five- or six-dimensional phase space for gyrokinetic or fully kinetic systems with three-dimensional real space), the PIC methods become more competitive. In five-dimensional (phase space) gyrokinetics for example, both Vlasov and PIC methods are widely in use. It should also be noted that there is numerical research into advanced grid techniques which may be able to avoid or mitigate the curse of dimensionality using, for example, sparse- or adaptive-grid methods. However these methods have yet to be successfully incorporated into large plasma 5D gyro- or 6D fully-kinetic codes.

3.2 Vlasov method

The first method is based on an Eulerian frame and focusses on the concept of sampling the distribution function on a fixed grid in phase space. It is frequently referred to as the Vlasov approach, which is a term that we shall use in this work. It uses $f(\mathbf{x}, \mathbf{v}, t)$ to represent the total amount of plasma density at time t at the phase space position

$(\mathbf{x}, \mathbf{v})$. It then uses a numerical scheme to evolve the distribution function on the grid using the fields, and to evolve the fields, which are also calculated on spatial grid, using the distribution function in combination with, in the case of the Vlasov-Poisson system, Poisson's equation.

Throughout this work, we perform Vlasov simulations using the MPI-parallelized nonlinear Vlasov-Poisson code viVLASOV1D, developed by M. Kraus using techniques described in [27, 28, 29]. The code uses a novel set of methods called variational integrators [30, 31]. The distinctive property of variational integrators is that the discretization is performed at the level of the Lagrangian, from which discrete equations of motion can be obtained via the discrete variational principle. As a consequence, the equations obtained are exact equations of motion of a discrete system which approximates a continuous system, rather than discrete equations approximating the equations of motion of the continuous system. As a result, the variational integrators preserve the conserved quantities of the continuous system, which is not usually the case for classical integrators. In particular, the simulations preserve the total particle number, linear momentum, and total energy to the order of the machine accuracy.

There is a long history of different numerical methods used, which have different convergence and conservational properties, but a particularly interesting result is that first shown in the results of Cheng and Knorr [32], who demonstrated how the operators can be split into an advection and an interpolation step, which is an important part of the semi-Lagrangian method. A more detailed explanation of the semi-Lagrangian method, and a discussion of many different numerical methods for the Vlasov-Poisson system is presented by Filbet and Sonnendrücker [33]

3.3   Particle-in-Cell method

The second method, particle-in-cell (PIC) [34], is based on a Lagrangian frame, and samples the plasma density as a series of macro-particles, which represent parcels of density, and follows the motion of these parcels. The particles themselves are assigned a weighting, which represents the amount of phase space density represented by the particle, and a position in phase space (position and velocity). During the time step, the particles are evolved (or pushed) using the equations of motion combined with the velocity of the particle and the forces due to the fields interpolated at the location of the particle. The fields are then calculated at the new time step by mapping the particles and their densities on to a grid, on which Poisson's equation is solved.

As the PIC method relies on a random sampling of the system by a limited number of macro-particles, the simulations are prone to noise and to statistical variance. Of all of the particle trajectories present in the physical system, only a limited number of them are present in the simulated system, each one representing a large number of real trajectories. However, in some marginal cases, or systems very close to a physical threshold, it is possible that some of these trajectories which are not sampled are the very same trajectories which may govern the physical behaviour of the system. As a result of the random initial sampling of phase space, PIC simulations are not deterministic[1], which helpfully allows a measure of the uncertainty of the result to be obtained simply by repeating a simulation multiple times. On the other hand, the non-deterministic property can make a comparison of the same case with increasing number of markers problematic. To try to solve this problem, we shall perform many of the PIC results

---

[1]In reality, the random numbers are calculated using a pseudo random number generator (PRNG), and therefore, if the same seed is provided, then the results are deterministic and repeatable.

multiple times when we are interested in making claims about the accuracy of the results, where feasible.

Throughout this work, we perform PIC simulations using the δf version of the MPI-parallelized nonlinear Vlasov-Poisson code PIC1D-PETSc (PIC1DP), developed by W. Deng and described in [35, 36].

3.4 Normalization

For the purposes of simulation, it is convenient to normalize certain quantities. All length scales are normalized to the Debye length ($\lambda_{\mathrm{D}} \equiv \sqrt{T/m}\,\omega_{\mathrm{p}}^{-1}$), all velocities are normalized to the thermal velocity ($v_{\mathrm{th}} \equiv \sqrt{T/m}$), where T is a so-called temperature, or a measure of the width of a Maxwellian distribution in velocity space. All frequencies are normalized to the plasma frequency, $\omega_{\mathrm{p}} \equiv v_{\mathrm{th}}/\lambda_{\mathrm{D}} = \sqrt{\frac{ne^2}{\epsilon_0 m}}$, and all times to the inverse plasma frequency (ω_{p}^{-1}). Charges are normalized to the absolute value of the electron charge, e (where $q_{\mathrm{e}} = -e$).

3.5 Mode initialization

The conventional method for initializing a mode in the Vlasov-Poisson system is via an initial density perturbation $n(x, t = 0) = n(1 + A_0 \cos{(kx)})$. As discussed in §2, the long term behaviour of the system is dominated by the pole of equation 2.12 with the largest imaginary part. However, just as Landau argued that for large values of t, the potential must be proportional to $e^{\gamma t}$, it is clearly not true that after very short times, that the initial perturbation will have relaxed into the linear eigenmode. In practice, we see that as the system parameters are varied to bring the system closer to marginal stability, the system takes longer until the eigenmode is reached. This itself poses several problems.

Firstly, if interested in a weakly unstable mode with wave number $k_0 \ll k_{\mathrm{peak}}$, where k_{peak} is the wave number corresponding to the maximum growth rate, then modes with wave numbers which are integer multiples of k_0 may have a higher growth rate than the mode under investigation. We know that the amplitude of the mode of interest is initialized to our starting amplitude ($A = A_0$) and we can assume that the modes with higher multiples of k_0 have an initial amplitude close to machine precision ($A \approx 10^{-17}$). Therefore, if our mode of interest takes longer to form the initial mode than the time it takes for the shorter wavelength mode to grow to a non-negligible amplitude, then we cannot claim to be able to measure the growth rate of our mode. If we are interested in simulating our mode until saturation or beyond, then we are further restricted by the shorter wavelength mode. One solution to this problem is to apply filters to only permit a certain mode or modes, however this neglects the evolution of harmonics of the wave. Another solution is to restart the simulation, and rescale or filter the amplitudes of the different modes when restarting. We have found that continually rescaling the wanted mode to the small amplitude A_0, whilst setting all other mode amplitudes to 0 can be a successful way of obtaining the system with the perturbation that of the eigenmode, without the presence of shorter wavelength modes.

Another related problem occurs when wanting to simulate a multi-mode system where the modes have matching amplitudes, which is particlarly relevant for the two-mode system with matching growth rates. Even if the modes are initialized to the same initial amplitudes, and in the long time limit form linear modes with the same growth rates, this does not guarantee that the amplitudes will match, as the amplitudes of the

modes once the eigenmodes have been established may still differ. In this case filtering is not able to solve the problem, however the rescaling method can.

3.5.1 Improved method for mode initialization

A more elegant solution to this problem, which saves computational effort compared with the previously mentioned restart and rescaling approach, and removes the possibility for numerical errors to accumulate, is to calculate the analytical mode structure f_1 from equation 2.9, which, provided that we can obtain the complex frequency ω, we can do by performing a Fourier transform in space and time, and gathering the terms allows us to write,

$$f_1(x, v, t) = k\frac{e}{m}\frac{\varphi(x, t) \cdot \partial f_0/\partial v}{\omega - kv} \tag{3.1}$$

where we can assume that $\varphi(x, t = 0) \sim \varphi_{\text{amp}} \cdot e^{i(kx + 2\pi R)}$, where $0 \leq R < 1$ is a phase factor, usual set to zero.

In the case of a growing mode, we can, instead of initializing the system via a density perturbation, initialize the system with a perturbation according to equation 3.1, where the perturbation can be multiplied by a scalar to obtain the desired A_0. The system then exhibits exponential growth from the start. Some care must be made when numerically solving the dispersion relation that the initial guess leads to a solution on branch with the largest imaginary part [37]. This method has been implemented in the initialization routines of both viVlasov1D and pic1dp.

3.6 Delta-f method

One further optimization that can be applied before implementing the equations into a numerical tool is the so-called delta-f (or δf) method [38, 39]. Using this method, the distribution function is split into a constant part (f_0) and a variable part (δf). As f_0 is known analytically, it is only necessary to dedicate numerical resources (grid points or marker particles) to represent the variable part of f, namely $\delta f(x, v, t)$. The PIC code that we use in this work, Pic1dp, is capable of running in both full-f and delta-f, the PIC delta-f formulation of the Vlasov-Poisson equations implemented is presented in [40].

Chapter 4

Results

4.1 Landau Damping

We begin by looking at the system of linear Landau damping. The system is initialized with a population of electrons whose velocity space distribution is Maxwellian,

$$f_0(v) = \sqrt{\frac{T}{2\pi}} \exp\left(\frac{-v^2}{2T}\right) \tag{4.1}$$

(with $T = 1$). The system is perturbed spatially such that

$$f(x, v, t = 0) = f_0(v)(1 + A_0 \cos(kx)) \tag{4.2}$$

Assuming that A_0 is small, then the system undergoes linear Landau damping, with the amplitude of the perturbation decaying over time with decay rate γ.

We perform simulations of Landau damping using the Vlasov method. Figure 4.1 shows one such example where $k = 0.5$. In this case, where $\mathbb{Re}\{\omega\} \approx 1.42$, the velocity distribution has a significant negative gradient at $v_{\text{res}} = \mathbb{Re}\{\omega\}/k \approx 2.8$, and therefore the damping rate is strong ($\gamma \approx -0.153$). Figure 4.2 shows a similar case, but with $k = 0.2$. In this case, where $\mathbb{Re}\{\omega\} \approx 1.06$, the velocity distribution's gradient at $v_{\text{res}} = \mathbb{Re}\{\omega\}/k \approx 5.3$ is relatively flat, which gives a significantly lower damping rate ($\gamma \approx -5.5 \cdot 10^{-5}$).

Figure 4.3 shows a comparison between the decay rates as calculated from simulations such as those shown in figures 4.2 and 4.1, and the decay rates calculated via the dispersion relation (equation 2.12).

4.2 Bump-on-Tail instability

4.2.1 Growth rate

A species of faster particles is added to the system via a new distribution function $f_0(v)$, shown in equation 4.3 and plotted in figure 4.4. The new population has a density of one tenth of the total density, $n_b = 0.1$; a velocity offset, $v_b = 4.5$; and a temperature $T_b = 0.25$, using the notion of temperature to denote the width of a Maxwellian distribution in velocity space. The total velocity space distribution function is therefore

$$f_0(v) = (1 - n_b)\sqrt{\frac{T_p}{2\pi}} \exp\left(\frac{-v^2}{2T_p}\right) + n_b\sqrt{\frac{T_b}{2\pi}} \exp\left(\frac{-(v - v_b)^2}{2T_b}\right) \tag{4.3}$$

17
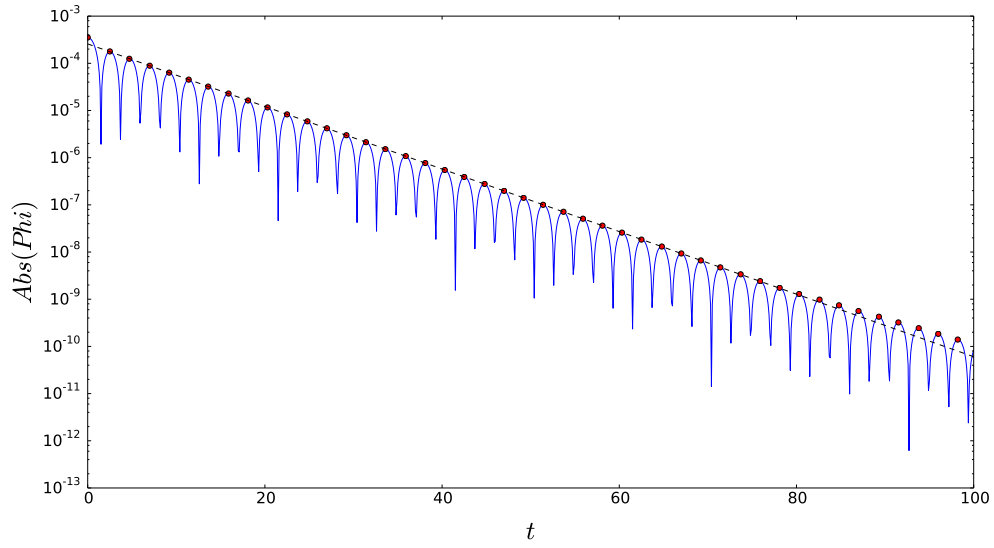
Figure 4.1: Strong Landau damping: Time evolution of $\phi(x = x_{\text{ref}}, t)$ for $k = 0.5$ showing how the decay rate (γ) is fitted — the dashed line, which has been fitted via the extrema of ϕ (the maxima of $|\phi|$), has a gradient of $\gamma_{\text{measured}} = -0.1527$. Simulation obtained using the Vlasov code.
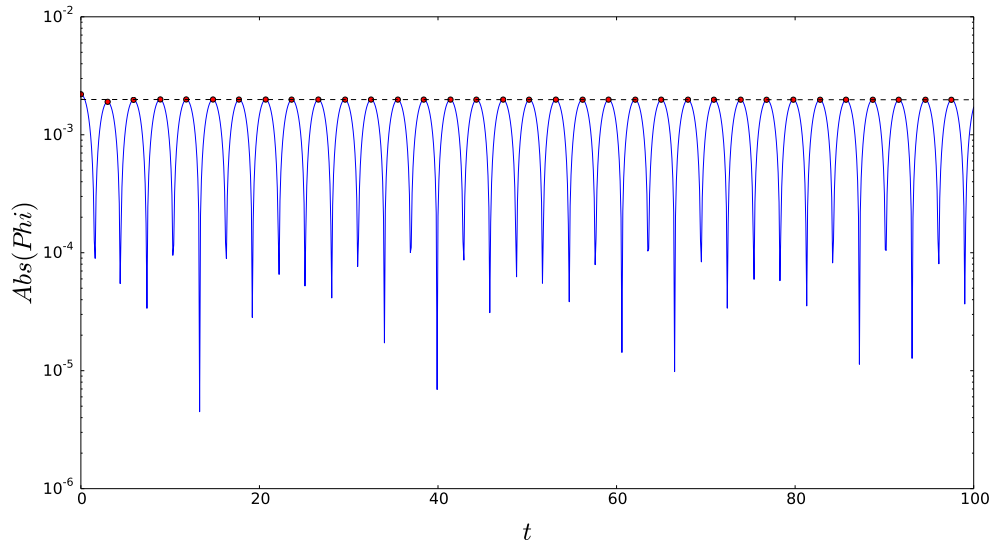
Figure 4.2: Weak Landau damping: Time evolution of $\phi(x = x_{\text{ref}}, t)$ for $k = 0.2$. $\gamma_{\text{measured}} = -5.26 \cdot 10^{-5}$.
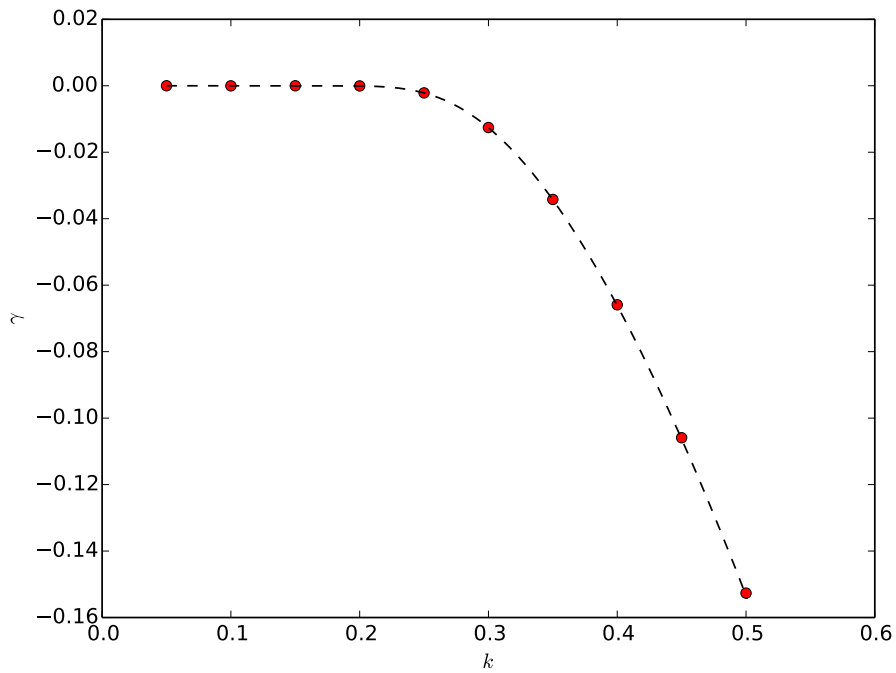
Figure 4.3: Landau damping decay rate (γ) as a function of the wave number k. Solid points correspond to numerical fits of Vlasov code simulations (including figures 4.2 and 4.1). The dashed line corresponds to the decay rates obtained via numerically solving the dispersion relation (equation 2.12).

(where $T_p = 1$) and is shown in figure 4.4. Note that as $n_b \to 0$, $f_{0,\text{bump-on-tail}}(v) \to f_{0,\text{Landau damping}}(v)$. As with the case of Landau damping, the system is initially excited by a small amplitude density perturbation. So just as in equation 4.2

$$f(x, v, t = 0) = f_0(v)(1 + \text{A}_0 \cos(kx))$$

Again, we begin by performing simulations using the Vlasov code. Figure 4.5 shows the electric potential sampled at a particular spatial reference point as a function of time, $\phi(x = x_{\text{ref}}, t)$. The high frequency oscillations corresponding to the real frequency can be seen, which makes separating the mode amplitude from this oscillation, as well as fitting the growth rate difficult.

By instead looking at a quantity which does not rely on sampling the potential at a single location, but rather accounts for all locations, then the high frequency oscillation can be removed, and the total magnitude evolution of the perturbation can better be seen. The square root of the electric field energy density, or electric field energy per unit length,

$$A(t) = \sqrt{\frac{1}{L} \int_0^L E(x, t)^2 \mathrm{d}x} \tag{4.4}$$

is chosen. By including the square root, the growth rate of this quantity is the same as the growth rate of the potential and of the density perturbation. For comparison with other articles which do not include the square root, the square of this quantity can be taken, and therefore also the growth rates doubled.

The same case as in figure 4.5 is presented in figure 4.6, but now showing the evolution of the new metric, A. It is clearly visible that there is an initial transient phase between $t = 0$ until $t \approx 25$. There is then a short linear period where the growth is purely exponential, before the saturation period, where the growth rate decreases leading to a maximum at $t \gtrsim 50$. In this example, the linear growth rate has been measured by performing a linear fit on the logarithm of the values within the designated fitting region ($30 < t < 40$). In contrast to figure 4.5, in figure 4.6, it is clear that the simulation is formed of an initial, or transient, phase; a linear growth phase; and a nonlinear saturation phase.

A similar procedure is performed for simulations with different k (and therefore different $L = 2\pi/k$), and the growth rates are shown in figure 4.7a. These values agree very well with the values calculated from the dispersion relation (shown with the dashed line). Additional points have been calculated using a similar method, but where, for convenience, the system was initiated with an initial condition directly corresponding to the linear growth phase as outlined in §3.5.1. These growth rates, plotting in figure 4.7b also agree well with the values from the dispersion relation, but can be calculated in significantly shorter time.

In figure 4.8, the transition towards the linear regime is shown. Figure 4.8d shows the phase space structure of the linear mode, with good agreement with the calculated mode, which is shown in figure 4.12a.

Convergence

For each of $k = 0.3$ and $k = 0.48$, short time length simulations have been carried out at a number of different resolutions. As can be seen in figure 4.7, $k = 0.3$ corresponds to one of the highest growth rates ($\gamma = 0.198$) of the system, whereas $k = 0.48$ is very close to marginal stability with a very small growth rate ($\gamma = 0.00520$) — from the dispersion relation, marginal stability occurs at $k = k_{\text{cr}} \approx 0.482$. In figure 4.9a, where
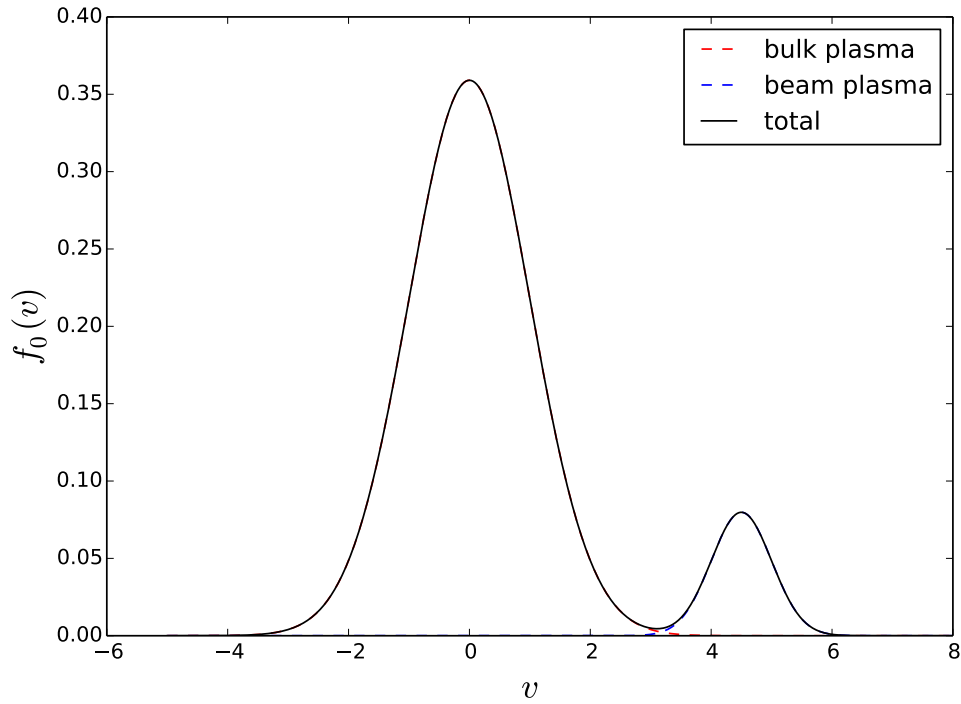
Figure 4.4: The velocity space distribution function for the bump-on-tail system ($n_b = 0.1$, $v_b = 4.5$, and $T_b = 0.25$).
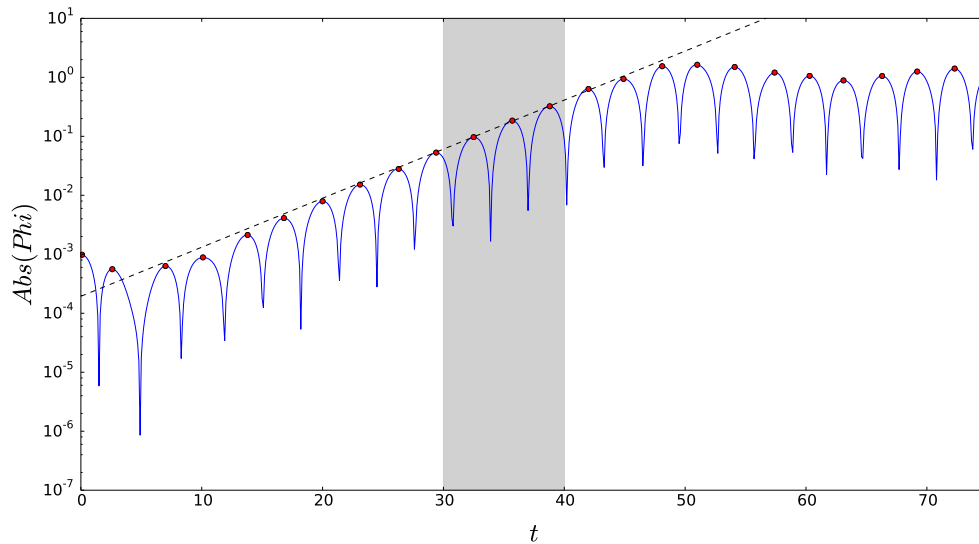

Figure 4.5: Growth due to the bump-on-tail instability. $\phi(x = x_{\mathrm{ref}}, t)$ as a function of t plotted for $k = 0.3$. The dashed line fits the growth rate via the maxima (marked with points) in the shaded region.
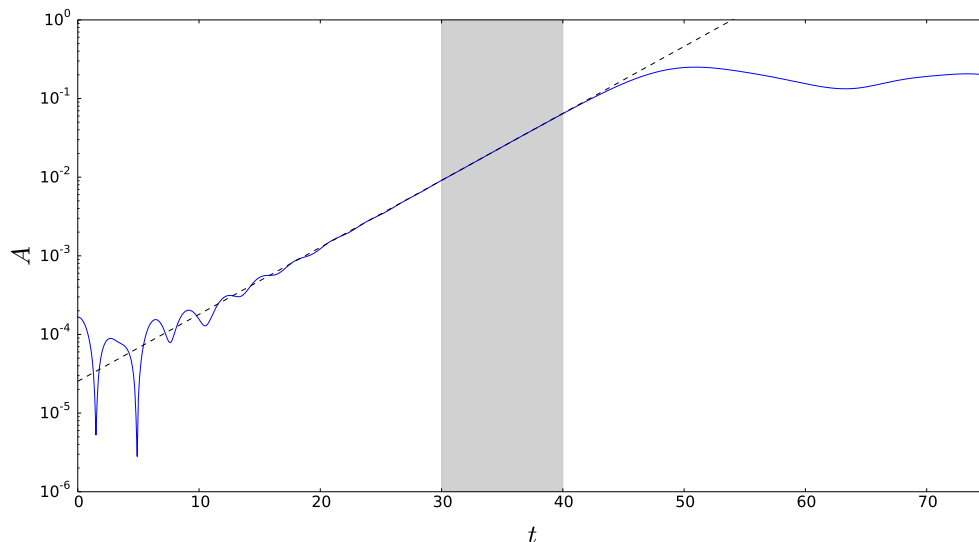
Figure 4.6: Growth due to the bump-on-tail instability. In contrast to figure 4.5, here the higher frequency oscillation of the wave is removed by plotting the square root of the total electric field energy density $A(t) = \sqrt{\frac{1}{L}\int_0^L E(x,t)^2 \mathrm{d}x}$ as a function t plotted for $k = 0.3$ (the same run as in figure 4.5). The dashed line fits the growth rate in the shaded (linear) region.

(a)   (b)
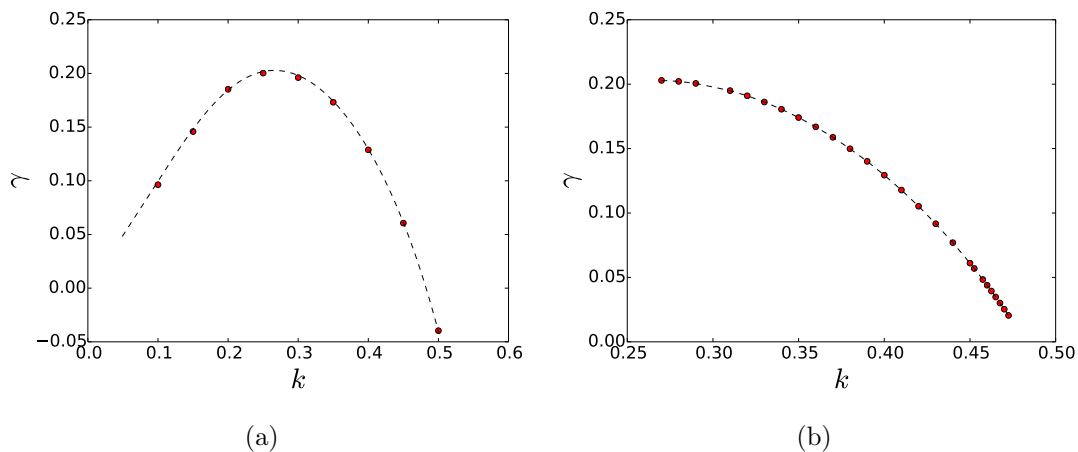
Figure 4.7: Bump-on-Tail growth rate (γ) as a function of the wave number k as measured using Vlasov code simulations. Lines correspond to the growth rates obtained from the numerical solution of the dispersion relation (equation 2.12). (a) Points obtained from simulations such as figure 4.6, initialized via a sinusoidal density perturbation. (b) Points obtained from simulations similar to those in (a), but in which the linear mode had been initialized from the start (as described in §3.5.1).

22

(a) $t = 0$

(b) $t = 5$

(c) $t = 15$

(d) $t = 30$
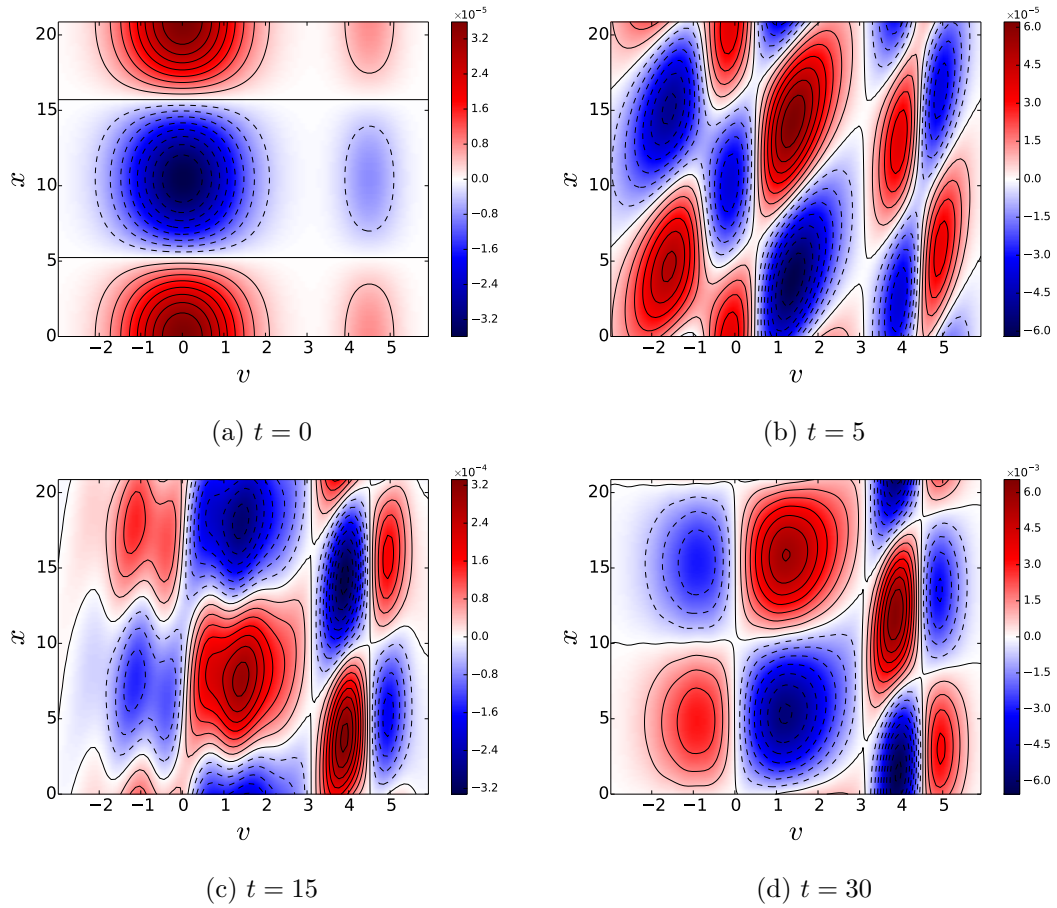
Figure 4.8: Snapshots of the mode structure $f_1(x, v) = f(x, v) - f_0(v)$ in the (x, v) plane for $k = 0.3$ — the same case as figures 4.5 and 4.6. The transient phase (b) and (c) stage occurs between initialization (a) and the establishment of the linear mode (d).

$k = 0.3$, the error $\left(\left| \frac{\gamma_s - \gamma_r}{\gamma_r} \right| \right.$, where γ_s is the growth rate obtained from the simulation , and γ_r is the growth rate obtained from the dispersion relation), quickly diminishes to approximately 1% for $(nx, nv) \geq (2^6, 2^5)$. However, in figure 4.9b, where $k = 0.48$, the error only approaches 1% for $(nx, nv) \geq (2^9, 2^{10})$.

We perform also the same cases as PIC simulations, but varying the number of markers and the resolution of the grid on which the fields are evaluated. These data are presented in figure 4.10. We observe no dependence of the error on the number of velocity space points in the grid, but once again, we see that the case of $k = 0.48$ is more demanding than that of $k = 0.3$. We see that an error of approximately 1% is reached for $k = 0.3$ at $np \approx 2^{10} \cdot 10^3$ for the worst case, and at $np \approx 2^7 \cdot 10^3$ for the median case. However, for $k = 0.48$ even $np = 2^{15} \cdot 10^3$ gives a median error estimate of between 1 and 10 %.

The explanation for this discrepancy can be seen in figure 4.11. There, the linear mode, $f_1(x, v)$, of the system is computed, and in both cases normalized such that the density,

$$n(x) = n_0(1 + 10^{-6} \cos{(kx)})$$

For each point v', the slice $f_1(v = v', x)$ is taken. For small amplitude, it is assumed that this slice is a pure sinusoidal function, with wave number k, therefore a Discrete Fourier Transform (DFT) is performed over x. The DFT's fourier components are then plotted in figure 4.11, where (a) and (b), or (c) and (d) uniquely define the mode structure.

In this figure, it is clear that the mode is smooth in v for $k = 0.3$. However, for $k = 0.48$, for $v \approx 3.1$, the mode has structure which is very narrow in v. And, as is shown in 4.11d, the phase, or angle of the mode also varies on a very narrow scale, which results in fine scale detail also in x.

A contour plot of $f_1(x, v)$ is presented in figure 4.12 for $k = 0.3$ and $k = 0.48$, as well as an intermediate case, $k = 0.44$. The intermediate case is presented as an aid to help show how the narrowing of the features takes place as the system tends towards marginal stability ($k \to k_{cr}$).

The physical reason for the narrowing of the phase space features is that as the system moves closer to marginal stability, the growth rate (γ), which is also the imaginary part of the complex frequency ($\omega_{complex} = \omega_{real} + i\gamma$), decreases. This not only decreases the drive, but also decreases the width of the resonance in the dispersion relation, see equation 2.12. It is this decrease in the width of the resonance in velocity space, which decreases the size of the phase space region in which the interaction takes place, which causes the narrowing of structures, and therefore is the source of the numerical resolution demand.

It should be noted that the more marginal case ($k = 0.48$) is of particular physical interest, with relevant growth rates for fast particle driven modes in ITER predicted to be in the range $\approx$ 1% [11, 12].

4.2.2 Mode saturation

As seen in figure 4.6 and discussed in §4.2.1, if the growth rate is positive, then the perturbation grows exponentially during the linear phase before saturating at a maximum. We denote the first maximum after the linear growth phase as the saturation amplitude A_{sat}, and the time at which this is reached as the saturation time, t_{sat}. For a series of simulations between $k = 0.27$ to $k = 0.4725$, figure 4.13 shows the saturation level of the energy density as a function of the wave number k.

Figure 4.14 shows the relationship between between the growth rate, γ, and the saturation amplitude, A_{sat}. Figure 4.14b, which plots the logarithms of both quantities,
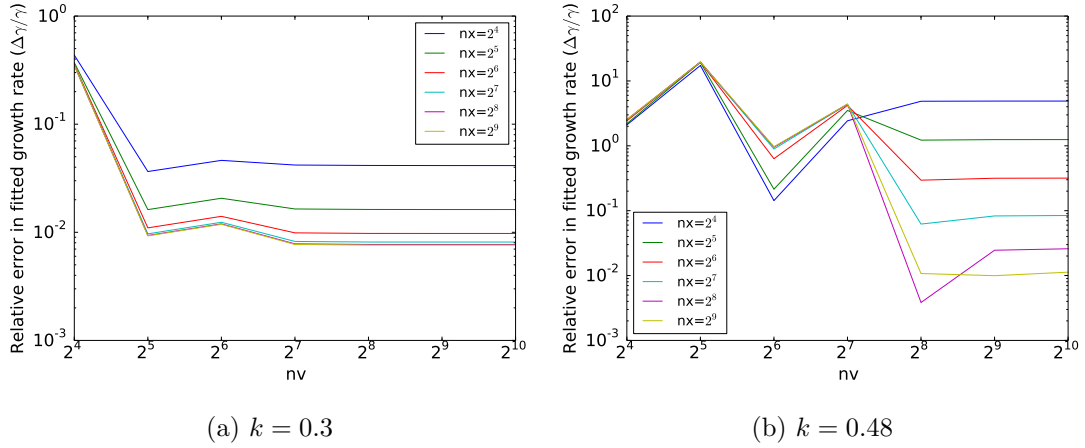
(a) $k = 0.3$ (b) $k = 0.48$

Figure 4.9: Resolution scans for $k = 0.3$ and $k = 0.48$ showing the relative error in the growth rate as a function of the simulation resolution. nx denotes the number of points between 0 and $L = 2\pi/k$ in real space. nv denotes the number of points in velocity space between $v_{\min} = -15$ and $v_{\max} = 15$.
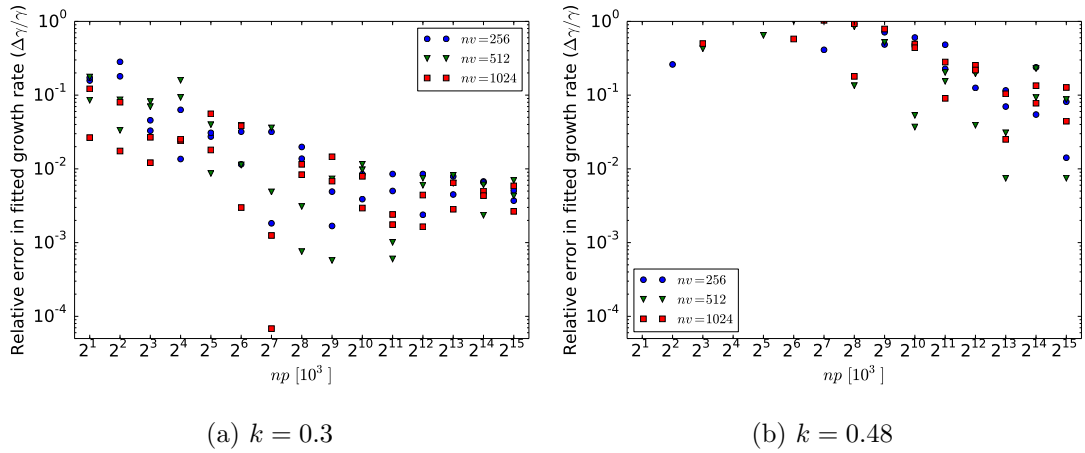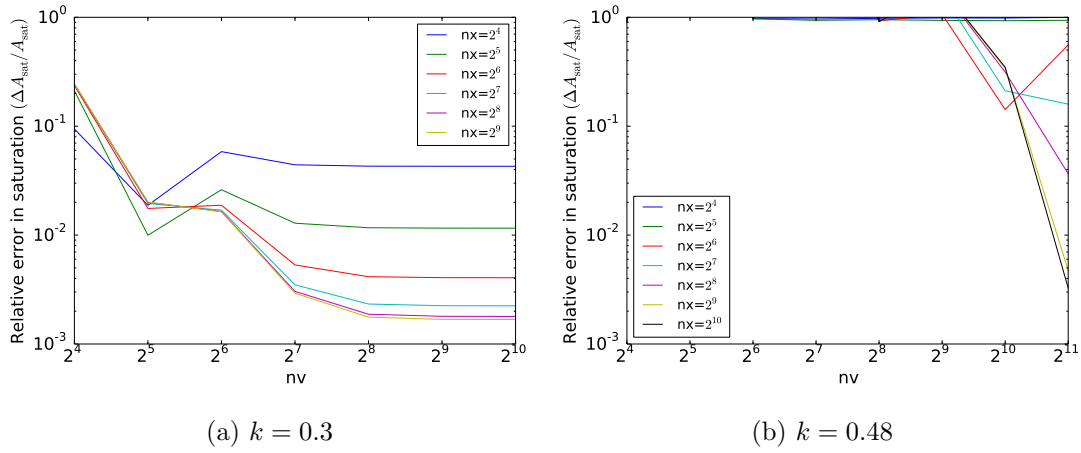

(a) $k = 0.3$ (b) $k = 0.48$

Figure 4.10: Resolution scans for PIC simulations with $k = 0.3$ and $k = 0.48$ showing the relative error in the growth rate as a function of the simulation resolution (number of markers). $nx = 1024$ denotes the number of points between 0 and $L = 2\pi/k$ in real space, and nv denotes the number of points in the velocity space between $v_{\min} = -15$ and $v_{\max} = 15$ of the grid on which the fields are calculated. Each simulation has been performed twice, because of, and showing, the inherent statistical uncertainty of the PIC method.
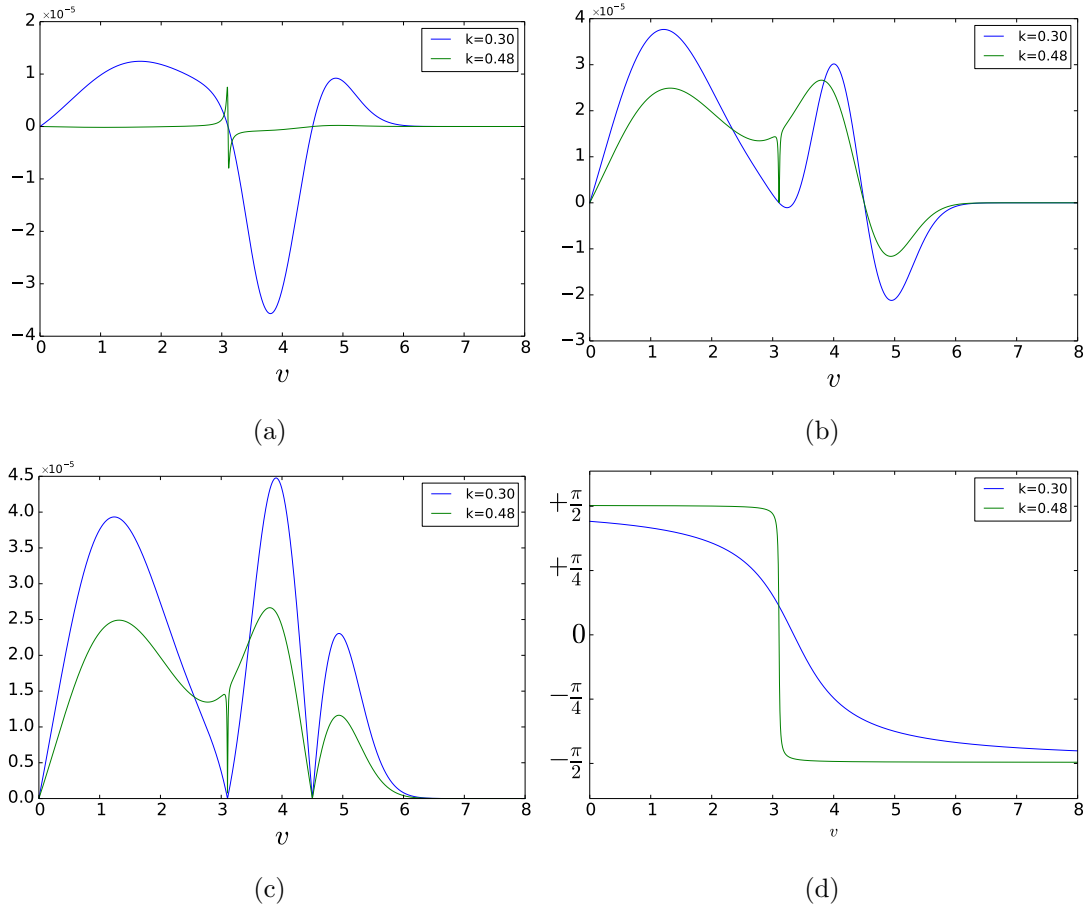
Figure 4.11: The mode structure for $k = 0.3$ and $k = 0.48$. The first fourier component of $f_1(x, v)$ is calculated over x at each point v.
(a) and (b): The real and imaginary parts of the mode structure plotted as functions of v.
(c) and (d): The magnitude and argument (with the jumps of $\pm\pi$ at the changes of sign in the real or imaginary parts removed for clarity) of the mode structure plotted as functions of v.
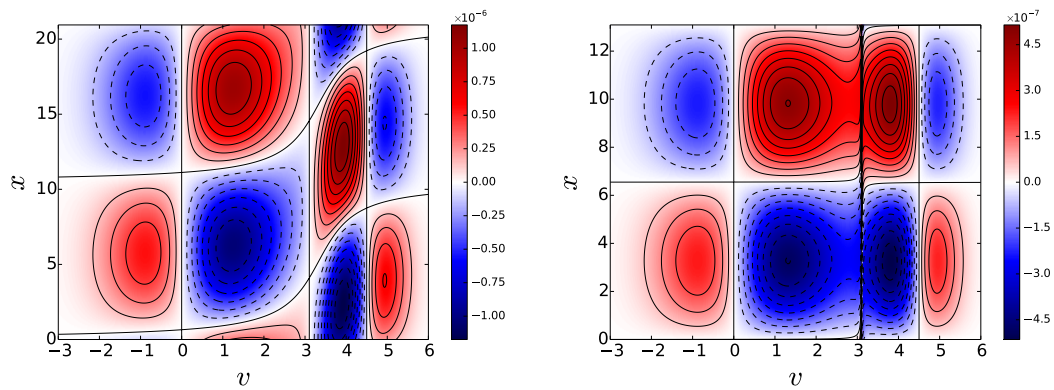
(a) $k = 0.3$

(b) $k = 0.48$
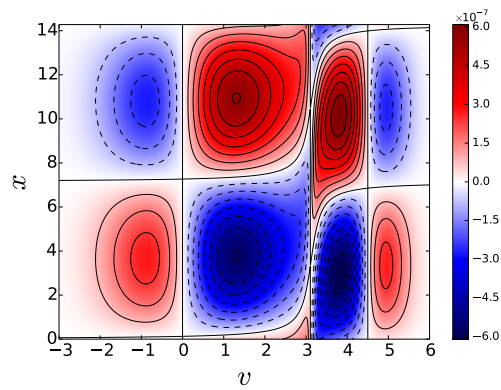
(c) $k = 0.44$

Figure 4.12: The mode structure of $f_1(x, v)$ in the (x, v) plane. (a) and (b) show the cases for $k = 0.3$ and $k = 0.48$ respectively, whilst (c) shows the intermediate case of $k = 0.44$.
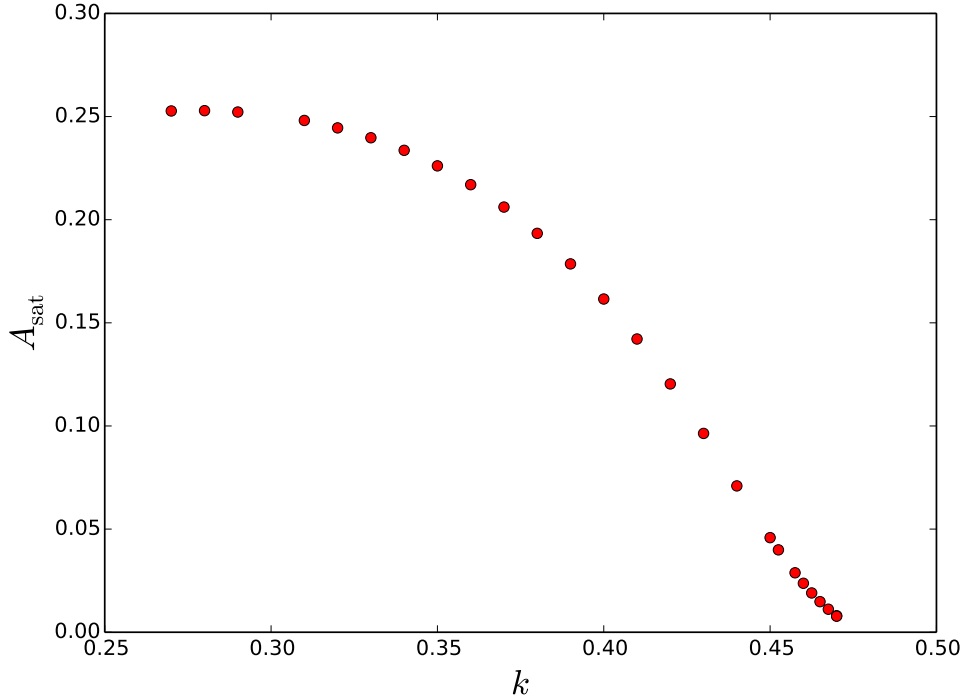
Figure 4.13: The value of the perturbation saturation amplitude (A_sat) as a function of k.

has a line with gradient 2 added which passes through the case of the lowest growth rate. It shows that, for small growth rates, the quadratic relationship between the growth rate and the saturation amplitude is nicely observed.

Convergence

See figure 4.15 for resolution scans comparable to figure 4.9, but for the saturation amplitudes (A_sat) rather than the growth rates (γ). As, unlike for the growth rate, there is no easily solvable analytical expression for the saturation amplitude, the "true" saturation level is obtained by running a simulation at the highest resolution, but with a smaller timestep ($\Delta t = 5 \cdot 10^{-2}$).

4.3 Multi-mode simulations

We next introduce systems with multiple perturbations. We choose to define a metric which allows us to see the relative perturbation amplitudes of multiple modes or harmonics. We therefore define $\widetilde{N}_k$ such that it corresponds to the amplitude of the density perturbation with wave number k. We note that, in the single mode cases discussed in §4.1 and §4.2.1, this is the quantity used to set the initial perturbation, $A_0 \equiv \widetilde{N}_{k_0}(t = 0)$. We also note the relationship between the density perturbation $\widetilde{N}_k$ and the square root energy density A in the case of a pure sinusoidal mode of wave number k,

$$\widetilde{A}_k = \frac{\widetilde{N}_k}{2k}$$

28

(a) Linear scales.

(b) Logarithmic scales.
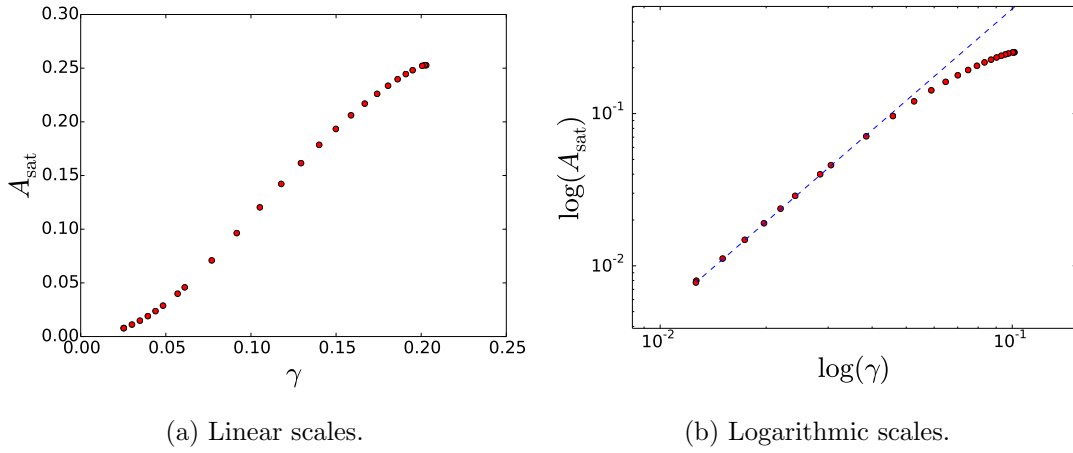
Figure 4.14: The value of the perturbation saturation amplitude (A_sat) as a function of the growth rate γ.

(a) $k = 0.3$

(b) $k = 0.48$

Figure 4.15: Resolution scans for $k = 0.3$ and $k = 0.48$ showing the relative error in the saturation amplitude as a function of the simulation resolution. nx denotes the number of points between 0 and $L = 2\pi/k$ in real space. nv denotes the number of points in velocity space between $v_\text{min} = -15$ and $v_\text{max} = 15$.
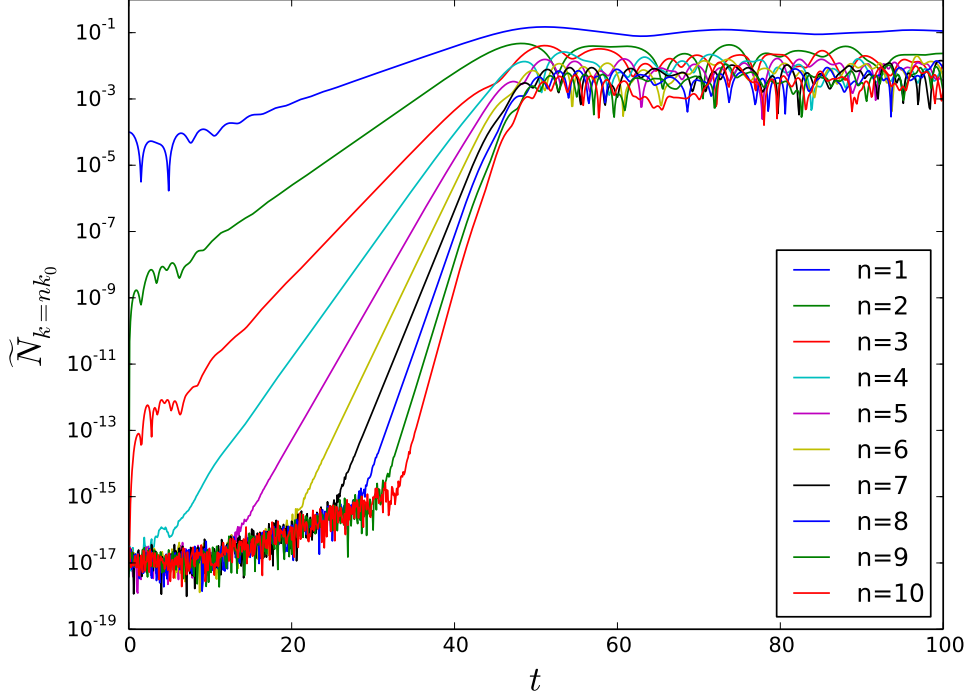
Figure 4.16: $\widetilde{N}_{n \cdot k_0}$ for $k_0 = 0.3$, $n = (1, \ldots, 10)$, $\widetilde{N}_{k_0}(t = 0) = 10^{-4}$. The same case as figure 4.5.

and the total square root energy density A in the case of multiple modes or harmonics,

$$A = \sqrt{\sum_k \widetilde{A}_k^2} = \sqrt{\sum_k \left(\frac{\widetilde{N}_k}{2k} \right)^2}$$

where $k = n \cdot k_0$ and n is the harmonic number such that the density,

$$n(x, t) = 1 + \sum_i \widetilde{N}_{k=k_i}(t) \cos \left(k_i x + P(t) \right)$$

where $P(t)$ specifies the phase of the perturbation.

Figure 4.16 shows the same single-mode case as shown in figures 4.5 and 4.6, but showing $\widetilde{N}_k(t)$ for $k = nk_0$, $k_0 = 0.3$ and $n = (1, \ldots, 10)$. This figure shows that the fundamental mode in the system dominates, but higher harmonics are also present with amplitudes, throughout the linear phase, of approximately

$$\widetilde{N}_{k=nk_0} = \left(\widetilde{N}_{k=k_0} \right)^n$$

This spectrum agrees with the findings for harmonics in [41].

For the multi-mode case, we again have the same $f_0(v)$ for the bump-on-tail system, but now we initialize the system as,

$$f(x, v, t = 0) = f_0(v) \left(1 + \sum_n \widetilde{N}_{n \cdot k_0}(t = 0) \cos(n \cdot k_0 x) \right)$$

In the limiting case, where $\widetilde{N}_{k_0}(t = 0) = A_0$, $\widetilde{N}_{nk_0}(t = 0) = 0$ for $n \neq 1$, we recover the initial condition for the single mode bump-on-tail system with initial amplitude A_0.

30

We choose k_0 such that the mode which we denote as n_0 corresponds to the mode with the maximum growth rate. We denote the number of modes present in the system as n_n, and these are chosen symmetically around n_0. Therefore, if n_n is odd, n_0 is an integer, and the modes n_0, $n_0 - 1$, $n_n + 1$, ... are present, centred around n_0, with the smallest and largest modes present corresponding to $n = n_0 \mp \frac{n_n - 1}{2}$. However, if n_n is even, then n_0 is a half-integer, and there is no "central" mode, the included modes lie between $n_0 \mp \frac{n_n - 1}{2}$. If n_0 is sufficiently large, it is expected that modes symmetrically around and sufficiently close to n_0 should have approximately the same growth rate. In the bump-on-tail system, we find that the maximum growth rate occurs when $k \approx 0.26$. We measure from the dispersion relation a value of $k_{\gamma_{\max}} = 0.266$, and from numerical simulations $k_{\gamma_{\max}} = 0.263$. We choose a convenient value which agrees to both of these values within 1%, $k_{\text{base}} = 0.264645$

In the case where all non-zero $\widetilde{N}_{nk_0}(t = 0)$ are equal, we denote this starting value as $\widetilde{N}_0$.

4.3.1   Growth rates via Vlasov simulation

For the following tests, we initialize the modes in the systems using a perturbation corresponding to the linear eigenfunction of the system as calculated via the linear dispersion relation, as outlined in §3.5.1.

One mode

We therefore introduce the system with $n_0 = 20$, and $n_n = 1$. We expect this case to be similar to the single mode cases in §4.2, except now there are multiple mode wavelengths in the periodic spatial domain, giving a reduced effective resolution in terms of grid points per wavelength. We therefore expect that the spatial resolution (nx) will have more stringent requirements.

Figure 4.17a shows a convergence scan for the growth rate γ. We conclude that for $n \approx 20$, $(nx, nv) \geq (2^9, 2^6)$ is required for an error in the growth rate close to 1%.

Three modes

We now keep $n_0 = 20$, but set $n_n = 3$. We look first at the growth rate of the mode which was present in the previous system $\gamma(k = 20 \cdot k_0)$. Figure 4.17b shows the error in the measured growth rate of this mode. Note that the error in the growth rate behaves very similarly to that in the case where $n_n = 1$ as a function of the resolution. We therefore claim that the presence of other modes in the system does not have a significant effect on the resolution required to resolve the growth rates.

More modes

Once again, we keep $n_0 = 20$, but increase n_n to 5. We now fix $nx = 1024$, and vary nv. Figure 4.18 shows the growth rates measured for the included modes in the range $n = 18 \ldots 22$, for each of $n_n = \{1, 3, 5\}$ run with resolutions $nv = 2^i$ for $i = 6 \ldots 10$. The black dash-dotted line shows the growth rates predicted from the linear dispersion relation.

This figure confirms that the linear growth rates are unaffected by the presence of additional modes in the system. It also agrees with our previous claim that increasing n_n does not increase the numerical resolution demands beyond those, previously discussed, in nx caused by the effective reduction in spacial resolution as a result of large mode numbers.
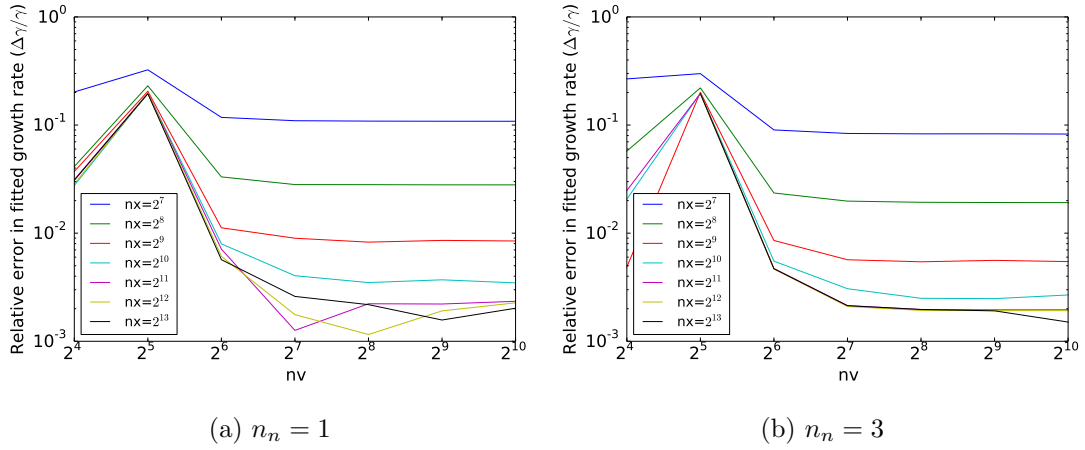
(a) $n_n = 1$ (b) $n_n = 3$

Figure 4.17: Resolution scans for $n_0 = 20$, $k_0 = 0.264645/n_0$ showing the relative error in the growth rate as a function of the simulation resolution. nx denotes the number of points between 0 and $L = 2\pi/k$ in real space. nv denotes the number of points in velocity space between $v_{\min} = -15$ and $v_{\max} = 15$.
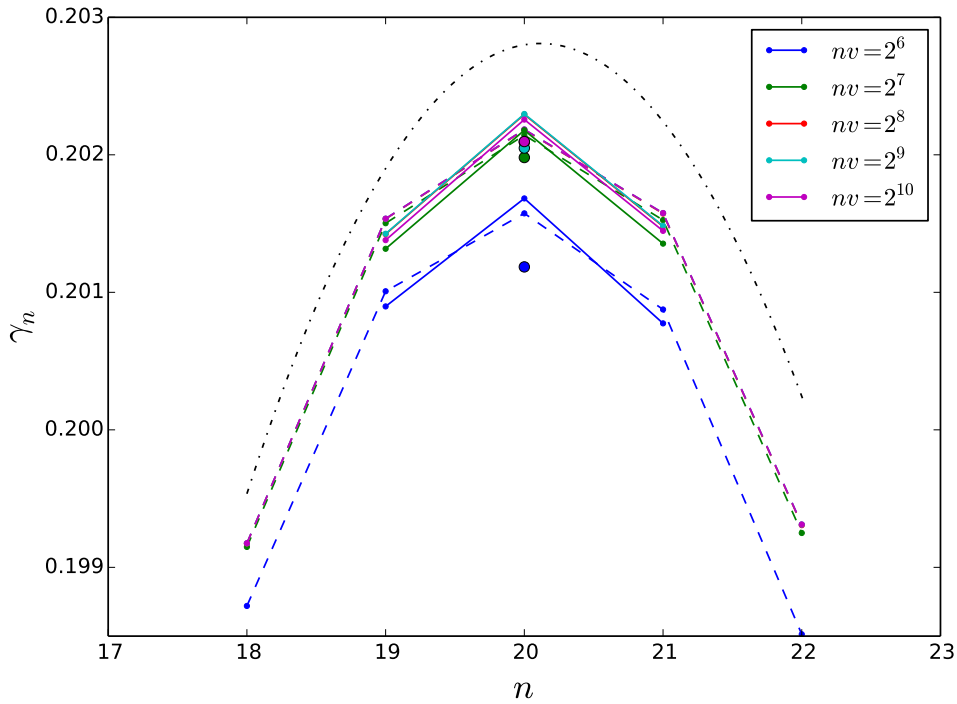

Figure 4.18: Growth rates for modes corresponding to $n = 18 \ldots 22$ for $n_0 = 20$, $k_0 = 0.264645/n_0$ simulated with $nx = 1024$ and varied nv. Large circles correspond to $n_n = 1$; solid lines with small points to $n_n = 3$, dashed lines with small points to $n_n = 5$; and the black dash-dotted line shows growth rates obtained from the dispersion relation.
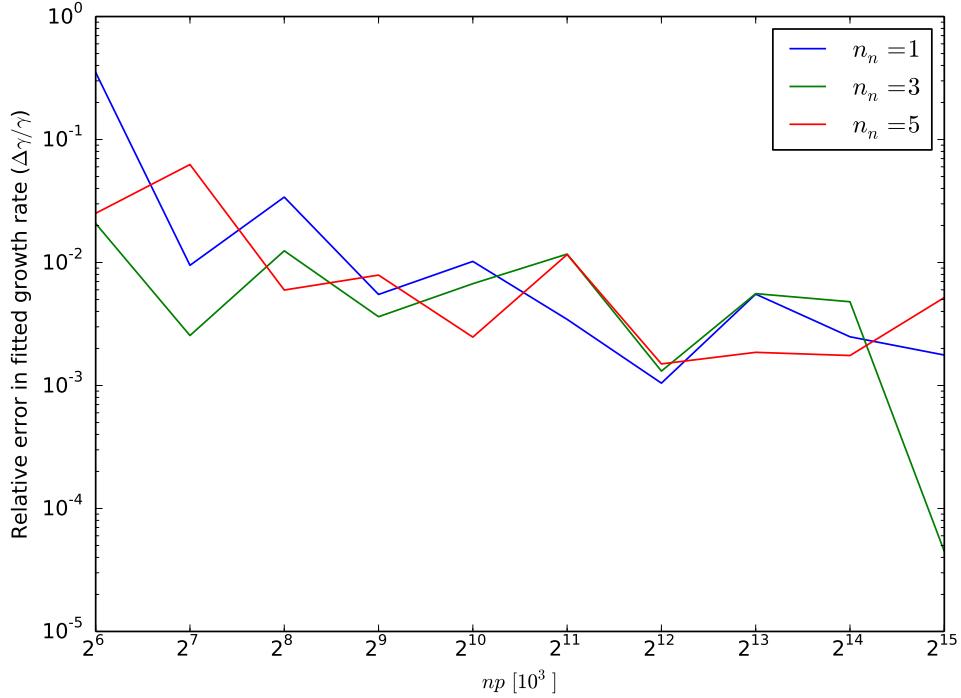
Figure 4.19: Relative error in the growth rate of the $n = 20$ mode, corresponding to $k = 0.264645$ as a function of number of particles (np), when different numbers of modes are present in the PIC simulations.

4.3.2 Growth rates via PIC simulation

We now repeat the same procedure, but perform PIC simulations. Again, we set $n_0 = 20$, and let $n_n = \{1, 3, 5\}$. The simulations were done with a grid of $nx = 1024$, $nv = 1024$, and we vary the number of particles (np) between $10^3 \cdot 2^{10}$ and $10^3 \cdot 2^{15}$. We find that, in order to obtain a good fit for the growth rate, a relatively long linear phase is required, and therefore the starting perturbation amplitude, $\widetilde{N}_0$, should be small, with $\widetilde{N}_0 = 10^{-6}$ allowing a clear fit for the growth rates.

By looking at just the growth rate of the central ($n = 20$) mode, just as in figure 4.17, we produce figure 4.19, which shows no dependence on the the number of modes in the system (n_n) and the rate of convergence of the growth rate with respect to the number of simulation markers.

Figure 4.20 shows the measured growth rates for the same systems as figure 4.18, the growth rates for all included modes between $n = 18$ and $n = 22$, for increasing numbers of simulation particles. We observe that almost all of the growth rates are within approximately 2% of the analytical curve, with most points with 1%, and that the statistical spread of the PIC method can be seen. We observe that several of the growth rates measured with $np \leq 2^{13} \cdot 10^3$ give larger error than those with higher numbers of particles, but that there is not a monotonically improving trend evident with single simulations with increasing np.
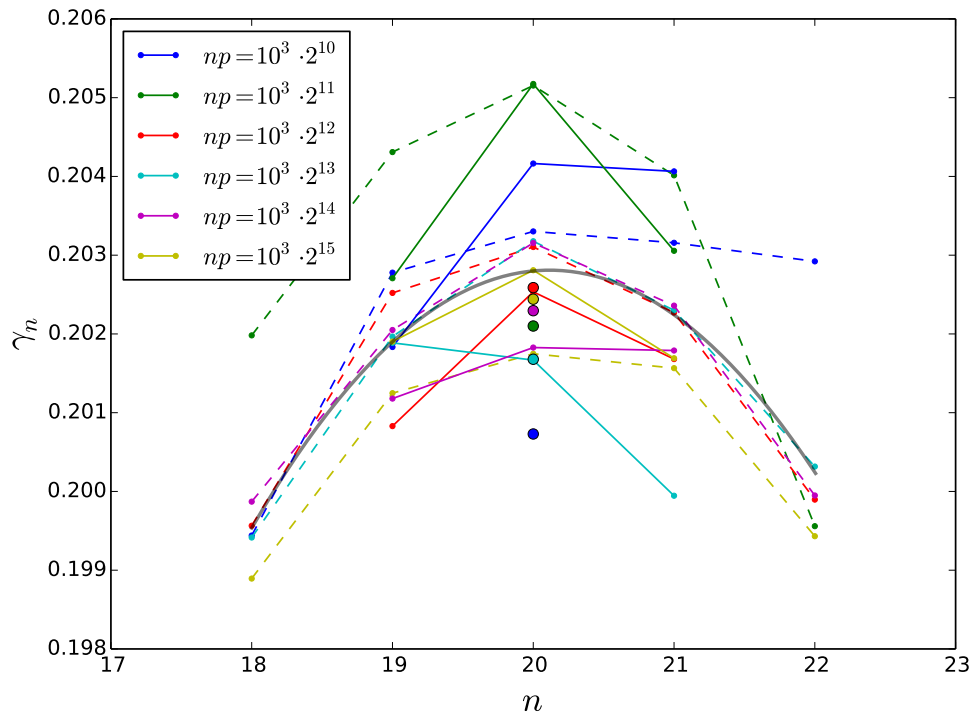
Figure 4.20: Growth rates for modes corresponding to $n = 18 \ldots 22$ for $n_0 = 20$, $k_0 = 0.264645/n_0$ simulated with $nx = 1024$ and varied nv. Large circles correspond to $n_n = 1$; solid lines with small points to $n_n = 3$, dashed lines with small points to $n_n = 5$; and the thick grey curve shows growth rates obtained from the dispersion relation.

34

4.4 Overlap of mode structures in phase space

It has previously been stated that realistic growth rates for future machines such as ITER may be of the order of $\gamma/\omega \approx 1\%$. It is therefore more realistic that there should be modes with much lower growth rates than were looked at in the previous section, which have much less overlap in their structures in phase space. We therefore investigate a range of systems with growth rates between $6 \cdot 10^{-3}$ and $1.5 \cdot 10^{-2}$, which, with ω between 1.05 and 1.03, gives a range of γ/ω spanning 1%. The growth rates are shown in figure 4.21.

With these parameters, a test system in which we can see the effect of resonance overlap is constructed. We widen the beam drive region by increasing the velocity spread of the beam to $T_b = 1.0$ and moving the beam further away from the bulk of the distribution ($v_b = 6.0$). We initialize two modes with mode numbers n_1 and $n_2 = n_1 + 1$. We choose $L = 2\pi/k_0$ such that the growth rates for mode numbers n_1 and n_2, (γ_{n_1} and γ_{n_2}) are equal. n_1 and n_2 are set to 6 and 7, chosen as this gives the mode spacings such that the islands are able to overlap with each other before overlapping with the boundary of the positive gradient region of velocity space (i.e. when $v_{7,\text{upper}} = v_{6,\text{lower}}$, $v_{7,\text{lower}} > \text{lowerbound}$ and $v_{6,\text{upper}} > \text{upperbound}$ (where $v_{n,(\text{lower,upper})} = v_{0,n} \pm \Delta v_n$, upperbound $= v_b = 6$, lowerbound $= v|_{f_{0,p} = f_{0,b}} \approx 4$))

We vary the growth rates, and therefore the island widths, by varying the beam density ($1 \cdot 10^{-3} < n_b < 2.5 \cdot 10^{-3}$), whilst keeping $\gamma_{n_1} = \gamma_{n_2}$ through small corrections to the wave numbers and to the box size. This corresponds to a range of the parameters needed to vary the island widths such that, when measured if simulated in isolation, they cross the threshhold of overlapping at saturation.

By using PIC simulations, we can look at particle trajectories to see how particles are trapped within the islands, which we can use to see the modes' width, and to see how particle transport is affected.

First we show the islands in isolated single-mode simulations. We do this by setting the box size to the wavelength of the mode in question. We measure the island width by empirically measuring the width of the the separatrix, by comparing the minimum position (in velocity space) of particles (density) which were initially above the resonant velocity to the maximum position of particles initially below the resonant velocity. Snapshots of phase space with particles coloured according to their initial velocities are shown in figures 4.23 and 4.25, and the evolution of the maxima is shown in figures 4.24 and 4.26 (in both cases for $n_b = 0.001$ and $n_b = 0.0024$ respectively). We note that in the $n_b = 0.001$ case, the islands are separated, yet in the case of $n_b = 0.0024$, the islands overlap. Figure 4.22 shows the evolution of the modes for $n_b = 0.001$. We see that the evolution of the mode amplitude is nearly identical for both modes, except for a minor difference in the initial conditions. Both modes saturate at amplitudes slightly in excess of $4 \cdot 10^{-4}$, namely approximately $4.25 \cdot 10^{-4}$ and $4.4 \cdot 10^{-4}$ for $n = 6$ and $n = 7$ respectively.

We now look at the above cases, but where both modes are present simultaneously. To this end, we increase the number of markers, and increase the box size such that 6 periods of the $n = 6$ mode and 7 periods of the $n = 7$ mode fit into the periodic box. Looking first at $n_b = 0.001$, we find that the amplitude evolution of the two modes, as seen in figure 4.27, are unaffected until after saturation, with the same saturation amplitudes and the same saturation time as in figure 4.22. Looking at the island position and width evolution, shown in figure 4.28, we see the same evolution as in the case of the single-mode runs (figure 4.24), and by looking at phase space in figure 4.29, we see the same structures, including the same widths and the same characteristic island wrapping.

Now, we look at $n_b = 0.0024$. This case differs from that of $n_b = 0.001$ in single-mode behaviour only in that it has a higher growth rates due to the larger gradients in the drive region as a result of the higher beam density. We observe, however, that the evolution of the mode amplitudes in the two-mode case, shown as the solid coloured lines in figure 4.30, differs substantially from that of the single-mode cases, which have been included as dashed grey lines. The multi-mode system exhibits the same linear behaviour, before an additional stage of growth occurs, which appears occur below the linear growth rate. The modes then continue to grow until they reach an amplitude of approximately 10^{-2}, this corresponds to a significant increase compared to the single-mode system levels, which give a saturation amplitude of only $2.4 \cdot 10^{-3}$.

By looking at the island position evolution in figure 4.31, we can make several observations. Firstly, the modes first overlap at a time $t \approx 600$, which is approximately the same time that mode amplitudes diverge from the single-mode amplitudes in figure 4.30. Secondly, we see that, once the modes have overlapped, particles are transported across the entire double island structure very quickly, such that by time $t = 750$, the minimum or maximum velocity space position of particles initially above or below one mode, and those above or below the other mode converge. Figure 4.32 shows a snapshot of phase space taken at $t = 650$, where the coloured regions are still identifiable, but have already become distorted. Figure 4.33 shows a snapshot of phase space at $t = 1300$, at the time of maximum island width, showing that particles with initial velocities between 4 and 6 are now almost completely mixed, and coherent structures within this range have been almost completely destroyed.

4.4.1 Distribution function flattening

Figure 4.34 shows the final velocity space distribution function from Vlasov simulations with $nx = 1024$, $nv = 2048$. We show three cases, one where the islands do not overlap ($n_b = 0.001$, figure 4.34a), one where the islands nearly overlap ($n_b = 0.00124$, figure 4.34b), and one where the islands clearly overlap ($n_b = 0.0024$, figure 4.34c). What we see in the case of the non-overlapping case is that two regions of the distribution function centred around the resonant velocities of the modes, exhibit local flattening of the distribution function, a flattening which reduces the drive of the modes, which is what causes the modes to saturate In the near-overlapping case, again we see flattening at the resonant velocities of the mode, but now the flattened region is wider. In comparison to the $n_b = 0.001$ case, the $\Delta f(v)$ line shows that there is a much narrower region between the structure of one mode and the structure of the other. If we look at the overlapping case, we see that the entire beam region has flattened, between $v = 4$ and $v = 7$, and all of the system's free energy has been released.

4.4.2 Sensitivity to initial conditions

In figure 4.35, we show three Vlasov simulations of the $n_b = 0.002$, $k_0 = 0.031478$ system. In figure 4.35c, both modes are initialized with amplitudes of $\widetilde{N}_0 = 10^{-6}$, and in figures 4.35a and 4.35b, modes 6 and 7 have their initial amplitudes increased to $\widetilde{N}_0 = 5 \cdot 10^{-6}$ respectively. It can be seen that the post-saturation evolution in all three cases is identical. Prior to the initially larger mode's saturation, both modes perform the same linear growth phase. Once the initially larger mode saturates, the lower mode, whose growth has already slowed as it begins to approach saturation, exhibits a phase of extra growth, which in turn enables the previous mode to grow again, until both modes reach a level much higher than their own original saturation levels. In the case of figure 4.35c, both modes follow the evolution of the initially lower mode.
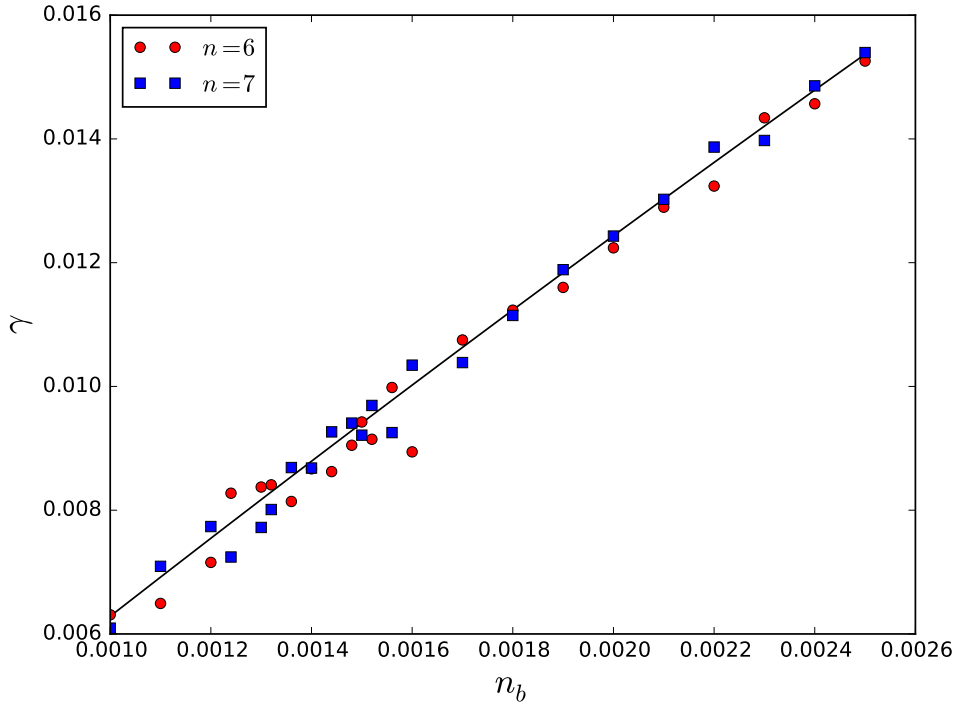
Figure 4.21: Growth rate (γ) for varying n_b. The black solid line corresponds to the linear growth rate measured from the dispersion relation, the red and blue solid points (circles and squares respectively) correspond to the growth rates of the 6th and 7th modes as measured in the linear growth phase of the PIC simulations in §4.4.
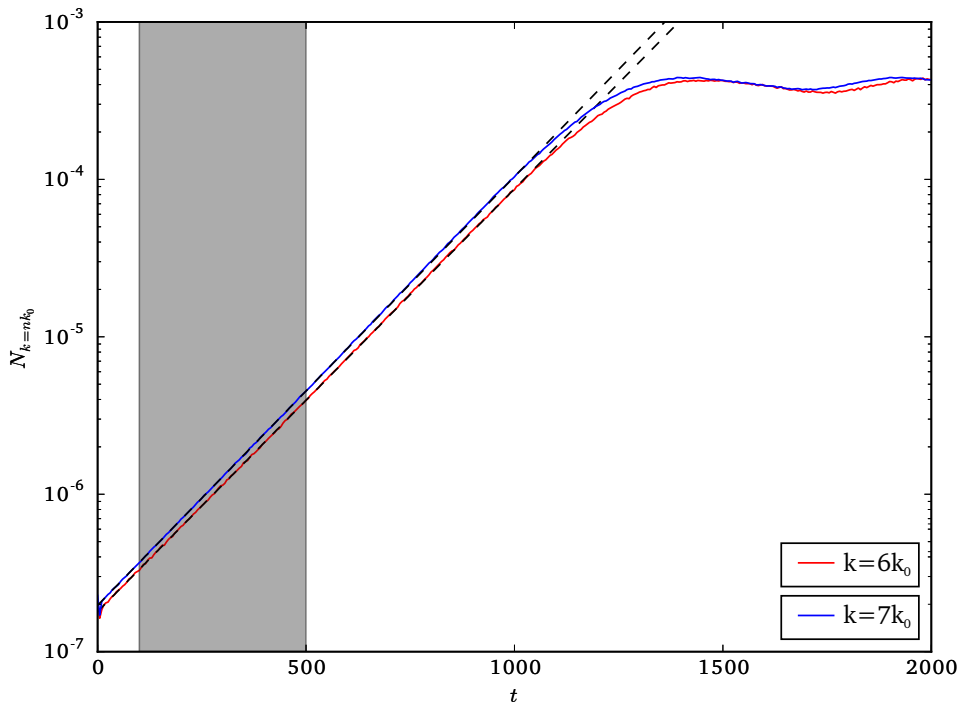

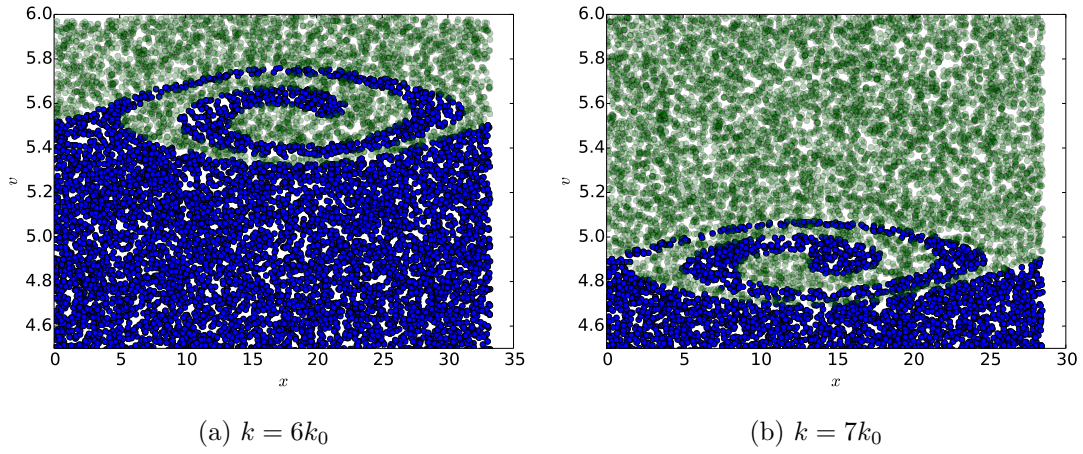Figure 4.22: Mode amplitude evolution for two one-mode simulations plotted on the same time axis. $n_b = 0.001$, $k_0 = 0.031521$.

(a) $k = 6k_0$

(b) $k = 7k_0$

Figure 4.23: A snapshot of diagnostic test particles' phase space positions close to the initial maximum width of the island, $t = 1500$. Particles are coloured based on their initial conditions: green particles began at $v > v_{res}$, and blue particles began at $v < v_{res}$. $v_{res} = \mathbb{Re}\{\omega\}/k$.

(a) $k = 6k_0$

(b) $k = 7k_0$
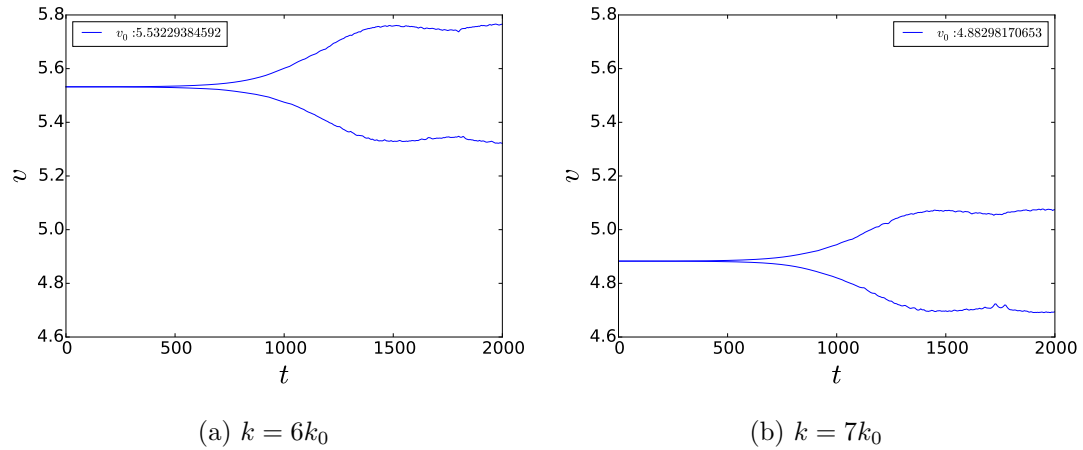
Figure 4.24: Evolution of the island position, calculated by taking the largest velocity space position of particles initially below the resonant velocity (i.e. blue particles in figure 4.23) and the lowest velocity space position of particles initially above the resonant velocity (i.e. green particles).

(a) $k = 6k_0$ (b) $k = 7k_0$
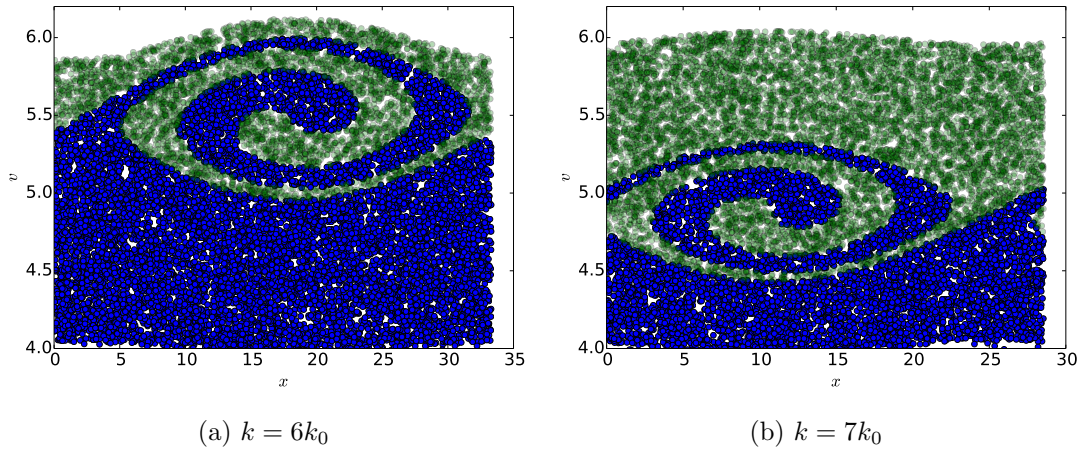
Figure 4.25: Island snapshots as in figure 4.23, but with $n_b = 0.0024$, taken close to the mode's initial maximum width, at $t = 750$.

(a) $k = 6k_0$ (b) $k = 7k_0$
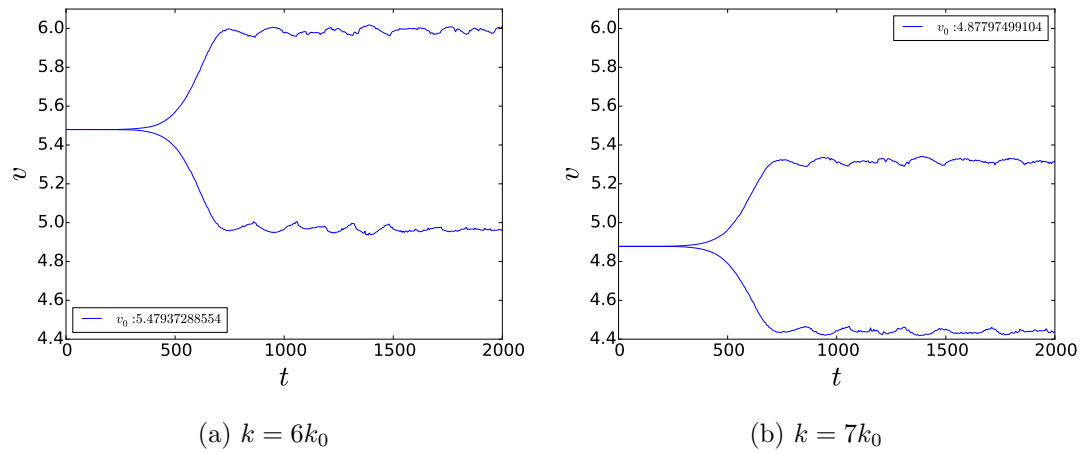
Figure 4.26: Evolution of the island positions, for $n_b = 0.0024$. Upper and lower bounds of the islands are calculated as in figure 4.24.
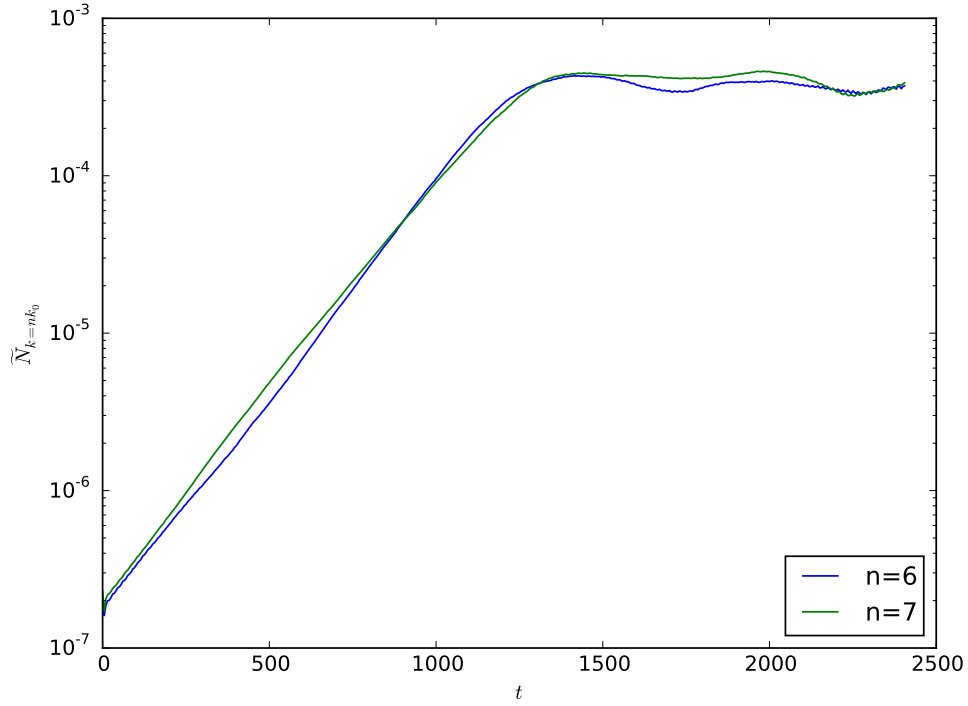
Figure 4.27: Evolution of the mode amplitudes in a two-mode simulation with $n_b = 0.001$. Parameters are the same as in figure 4.22, but with both modes present in the same simulation.
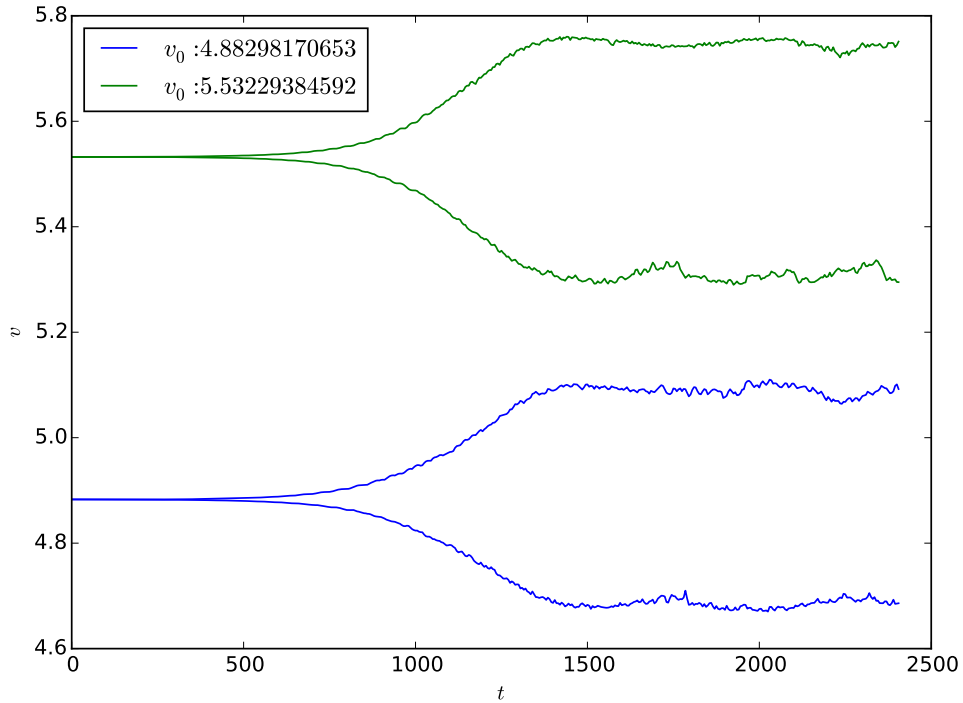
Figure 4.28: Evolution of the island positions for $n_b = 0.001$. This figure agrees closely with figure 4.24, which is the equivalent plot, but for the single-mode simulations.
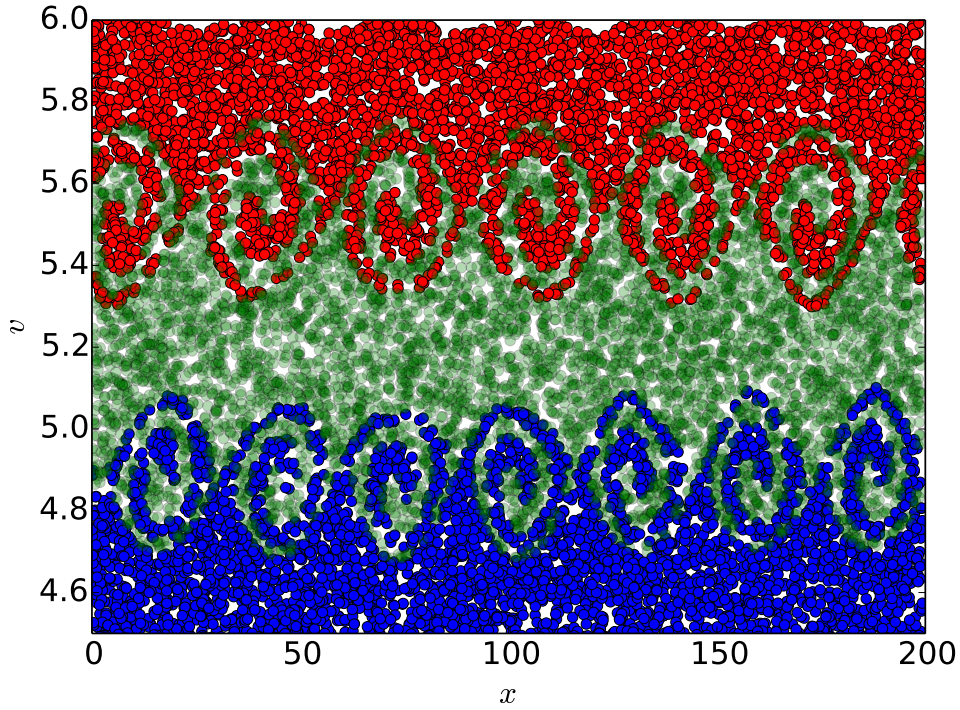
Figure 4.29: Island snapshot of $n_b = 0.001$ at time $t = 1500$. Blue particles have an initial velocity $v(t = 0) < v_{0,7}$, red with initial velocity $v(t = 0) > v_{0,6}$ and green with initial velocities in between, $v_{0,7} < v(t = 0) < v_{0,6}$
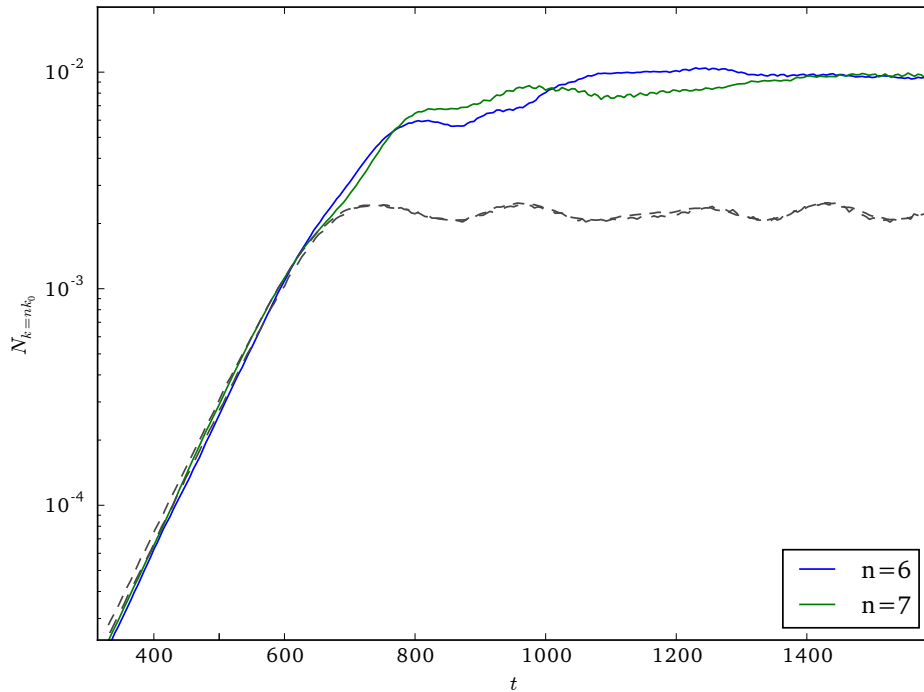
Figure 4.30: Evolution of the mode amplitudes for $n_b = 0.0024$ showing the amplitudes in single-mode cases in dashed grey lines, and in the case of the multi-mode case in the solid coloured lines.
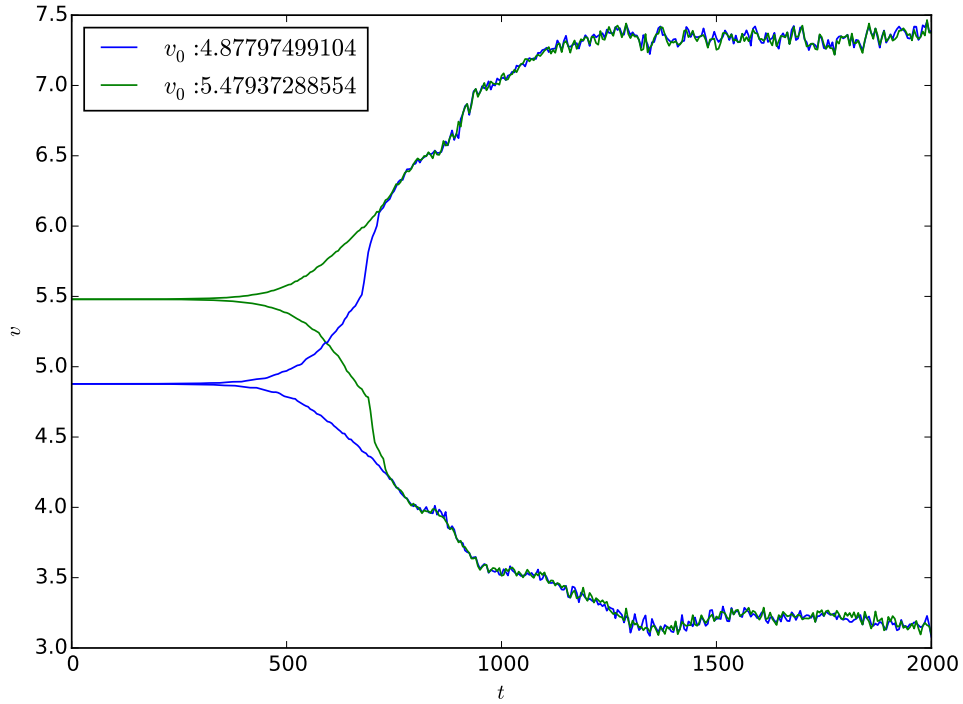
Figure 4.31: Evolution of the island positions for $n_b = 0.0024$. This figure contrasts strongly with figure 4.26, which is the equivalent plot, but for the single-mode simulations.
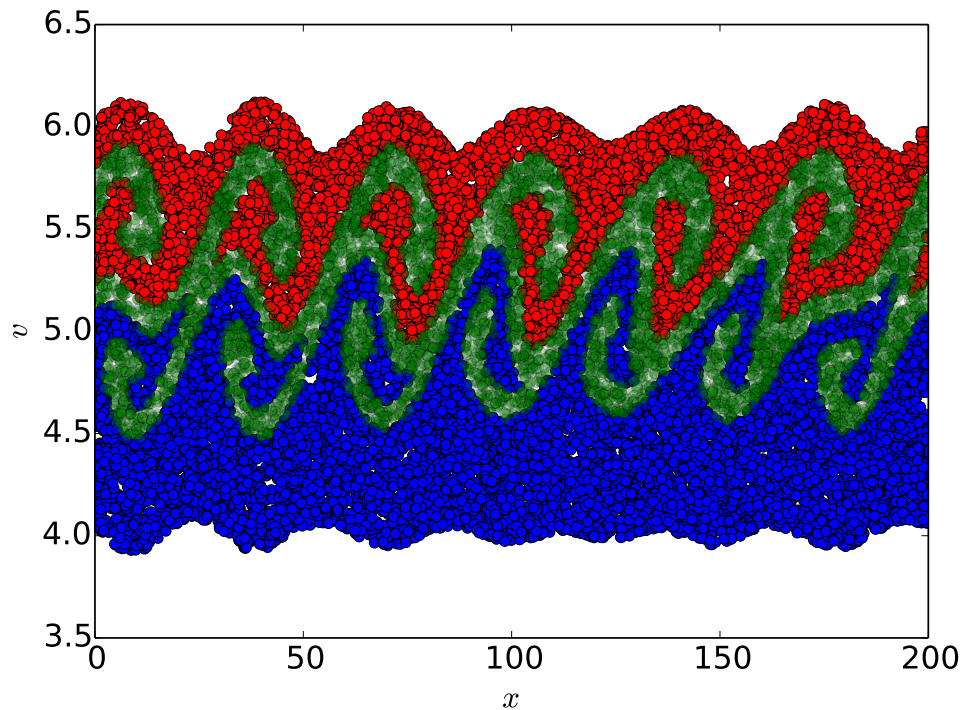
Figure 4.32: Island snapshot of $n_b = 0.0024$ at time $t = 650$. Blue particles have an initial velocity $v(t = 0) < v_{0,7}$, red with initial velocity $v(t = 0) > v_{0,6}$ and green with initial velocities in between, $v_{0,7} < v(t = 0) < v_{0,6}$

42
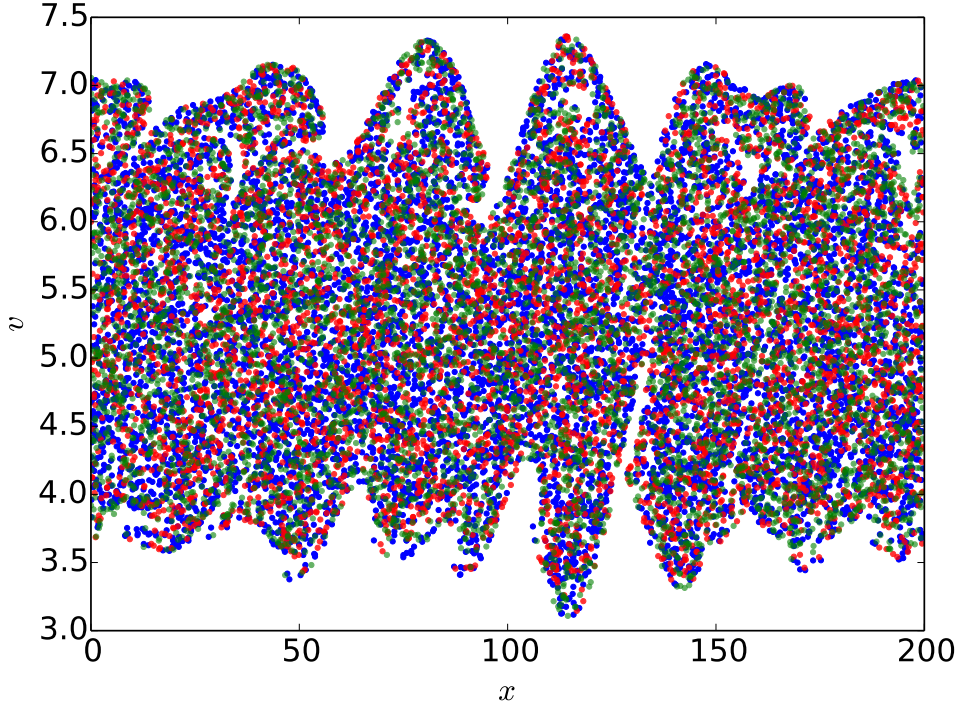
Figure 4.33: Island snapshot of $n_b = 0.0024$ at time $t = 1300$. Blue particles have an initial velocity $v(t = 0) < v_{0,7}$, red with initial velocity $v(t = 0) > v_{0,6}$ and green with initial velocities in between, $v_{0,7} < v(t = 0) < v_{0,6}$

This behaviour agrees with that observed in the case of overlapping toroidicity-induced Alfvén eigenmodes and reversed shear Alfvén eigenmodes [20].

4.4.3   Cross-island transport

Noting the qualitative differences in behaviours in the two examples presented above, we look for the transition between these two regimes. We choose to look at the amount of phase space density which has been transported across a point which lies between the two islands, which we define as $f_t(t)$, and we choose the point as $v' = 5.1$

$$f_t(t) = \int_{-\infty}^{v'} f(v, t) \, dv - \int_{-\infty}^{v'} f(v, t = 0) \, dv \tag{4.5}$$

This quantity is then normalized to the total density of phase space. The evolution of this quantity is plotted for $n_b = 0.001$ in figure 4.36a and for $n_b = 0.0024$ in figure 4.36b.

The dependence of the saturation level $f_t(t = \infty)$ of this quantity on the total beam density (n_b) is shown in figure 4.37. This step, or threshhold, corresponds to the critical beam density, $n_{b,c} = 0.0012$ at which the modes would not overlap in the case of a single-mode system. The high n_b linear trend can be explained because the beam has a higher total density, linearly proportional to n_b. Therefore, if redistribution is occuring, even without a qualitatively different behaviour, one expects the transported distribution f_t to increase with increasing n_b.
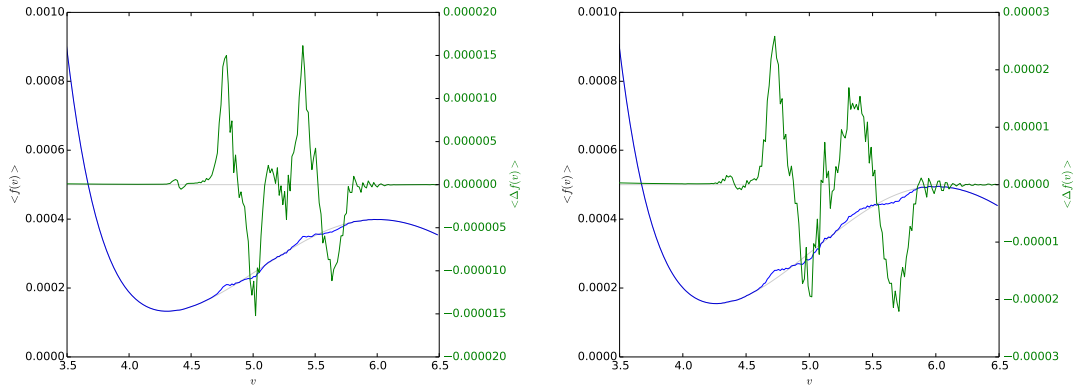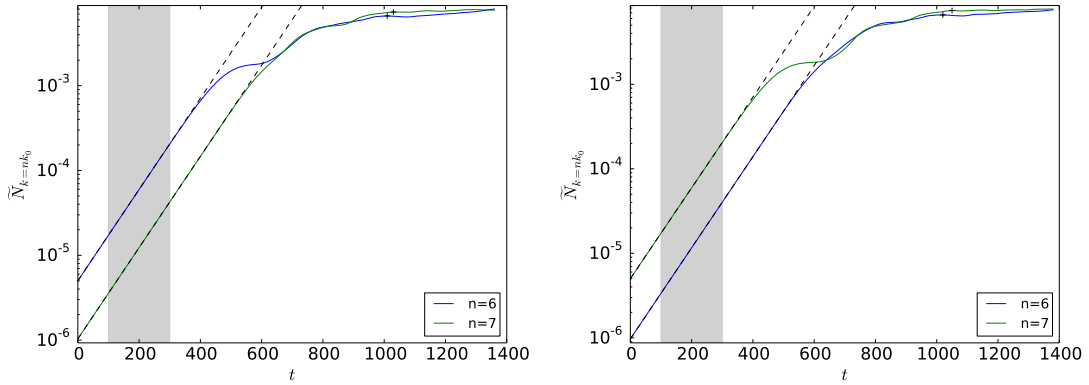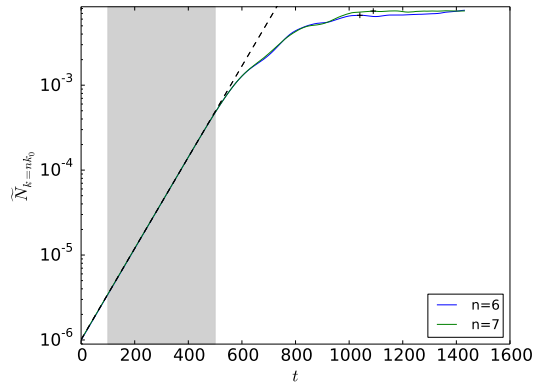
(a) $n_b = 0.001$

(b) $n_b = 0.00124$
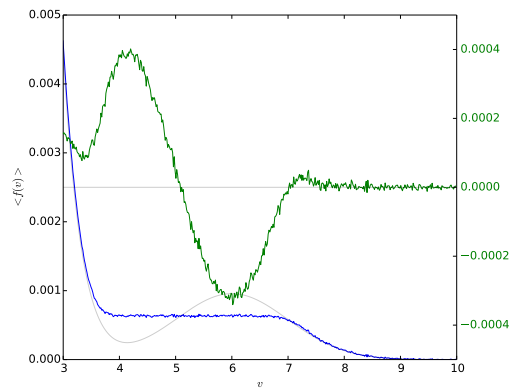
(c) $n_b = 0.0024$

Figure 4.34: $f(v)$ (blue, left axis) and $\Delta f(v)$ (green, right axis) plotted at the end of Vlasov simulations. The grey lines show the initial values of the distribution function.

(a) $\widetilde{N}_{\{6,7\}\cdot k_0}(t=0) = \{5,1\} \cdot 10^{-6}$

(b) $\widetilde{N}_{\{6,7\}\cdot k_0}(t=0) = \{1,5\} \cdot 10^{-6}$

(c) $\widetilde{N}_{\{6,7\}\cdot k_0}(t=0) = \{1,1\} \cdot 10^{-6}$

Figure 4.35: Mode amplitude evolution for $n=6$ and $n=7$, for three different sets of initial amplitudes.

Convergence

To study the relevant sensitivities and difficulties involved in these results, we select a limited set of cases, using figure 4.37 as a guide. We seek a point from the lower, non-overlapping region, a point from the higher, clearly overlapping region, and a point from the transition region, the marginally overlapping region. We choose $n_b = \{1.0, 1.24, 2.4\} \cdot 10^{-3}$ for these three cases.

We begin by focussing our attention to the case where $n_b = 1 \cdot 10^{-3}$, which was the non-overlapping case from §4.4, and the evolution of the islands was plotted in figure 4.28. By considering the evolution of the transported phase space density, $f_t(t)$, which has already been plotted for a single simulation in figure 4.36a. Now, we look at the sensitivity of this evolution by running PIC simulations with different numbers of markers (np), and running Vlasov simulations with different velocity space grid resolution (nv). This is shown in figure 4.38. We see that both the Vlasov (for $nv \geq 1024$) and PIC results agree, with a final saturated value of approximately $1 \cdot 10^{-7}$.

Next, we look at the case where $n_b = 2.4 \cdot 10^{-3}$, which was the overlapping case from §4.4, shown already in figure 4.31 (islands), and in figure 4.36b (transport). Again, we perform this with PIC and Vlasov simulations, varying the number of markers and velocity space grid points respectively. This is shown in figure 4.39. We see very good agreement, with all resolutions of PIC and Vlasov showing a saturated value of $3.6 \cdot 10^{-4}$.

Finally, we look at the intermediate case of $n_b = 1.24 \cdot 10^{-3}$, close to the transition seen in figure 4.37. Figure 4.40 shows the evolution of the transported phase space density ($f_t(t)$), showing different nv for Vlasov simulations in 4.40a, and different np for PIC simulations with $nv = 512$ in 4.40b. Here, we observe very interesting findings. Firstly, for an insufficient number of velocity space grid points ($nv = 256$ and $nv = 512$), the saturation level of the Vlasov simulations are approximately factor 10^2 larger ($f_t(t = 4000) \approx 10^{-4}$) than for higher resolutions simulations ($f_t(t = 4000) \approx 10^{-6}$). It should, however, be noted that for $nv = 256$, the linear growth rates were not correctly resolved (see table 4.1).

On the other hand, for lower values of np in PIC simulations in figure 4.40b, we observe that the majority of the simulations produce the lower value observed in the Vlasov case, with only one of five full length $np = 2^{14} \cdot 10^3$ simulations reaching the higher level of $f_t \approx 10^{-4}$. However, by doubling the number of markers to $np = 2^{15} \cdot 10^3$, this now becomes four from five. We attribute this to an increase in the probability of attaining the higher value with increased np, and recall from the description of the PIC method from §3.3 that the method has inherent noise, and statistical variance. As we further increase np up to a maximum of $np = 2^{19} \cdot 10^3$, we observe that most of the cases yield the higher value of $f_t \approx 10^{-4}$.

We look for any effect of the value of nv by also performing PIC simulations with $nv = 1024$ and $nv = 2048$ for broad ranges of np. These results are shown in figure 4.40c and 4.40d respectively. We do not observe a qualitative shift in the behaviour of the cross-island transport by increasing nv for PIC simulations.

Table 4.1: Growth rates for modes $n = 6$ and $n = 7$ for Vlasov simulations for three values of n_b. Analytical growth rates from the solution of the dispersion relation are shown in the table captions ($\gamma_{\mathrm{n=6}} = \gamma_{\mathrm{n=7}} = \gamma$).

(a) $n_b = 0.001$, $\gamma = 6.286 \cdot 10^{-3}$

nv	$\gamma_{n=6}$	$\gamma_{n=7}$
256	$2.66\cdot10^{-3}$	$8.30\cdot10^{-3}$
512	$6.29\cdot10^{-3}$	$6.11\cdot10^{-3}$
1,024	$6.33\cdot10^{-3}$	$6.33\cdot10^{-3}$
2,048	$6.29\cdot10^{-3}$	$6.28\cdot10^{-3}$
4,096	$6.28\cdot10^{-3}$	$6.26\cdot10^{-3}$
8,192	$6.28\cdot10^{-3}$	$6.26\cdot10^{-3}$
16,384	$6.28\cdot10^{-3}$	$6.26\cdot10^{-3}$

(b) $n_b = 0.0024$, $\gamma = 1.479 \cdot 10^{-2}$

nv	$\gamma_{n=6}$	$\gamma_{n=7}$
256	$1.50\cdot10^{-2}$	$1.52\cdot10^{-2}$
512	$1.46\cdot10^{-2}$	$1.46\cdot10^{-2}$
1,024	$1.46\cdot10^{-2}$	$1.46\cdot10^{-2}$
2,048	$1.46\cdot10^{-2}$	$1.46\cdot10^{-2}$
4,096	$1.46\cdot10^{-2}$	$1.46\cdot10^{-2}$
8,192	$1.46\cdot10^{-2}$	$1.46\cdot10^{-2}$
16,384	$1.46\cdot10^{-2}$	$1.46\cdot10^{-2}$

(c) $n_b = 0.00124$, $\gamma = 7.797 \cdot 10^{-3}$

nv	$\gamma_{n=6}$	$\gamma_{n=7}$
256	$6.43\cdot10^{-3}$	$9.52\cdot10^{-3}$
512	$7.60\cdot10^{-3}$	$7.55\cdot10^{-3}$
1,024	$7.99\cdot10^{-3}$	$7.99\cdot10^{-3}$
2,048	$7.78\cdot10^{-3}$	$7.77\cdot10^{-3}$
4,096	$7.80\cdot10^{-3}$	$7.78\cdot10^{-3}$
8,192	$7.79\cdot10^{-3}$	$7.78\cdot10^{-3}$
16,384	$7.81\cdot10^{-3}$	$7.79\cdot10^{-3}$
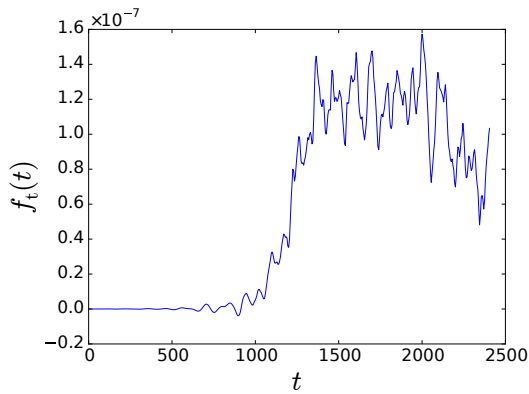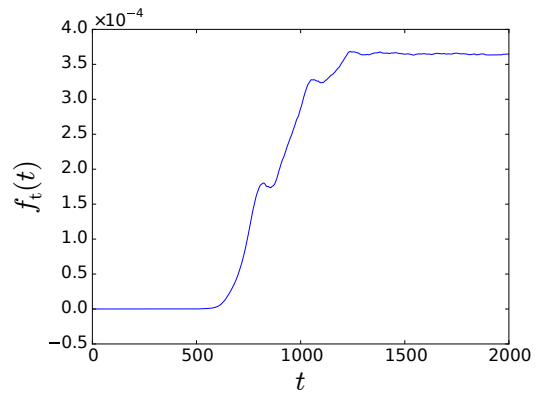
(a) $n_b = 0.001$ (b) $n_b = 0.0024$

Figure 4.36: Evolution of $f_{\mathrm{t}}(t)$ for $n_b = 0.001$ and $n_b = 0.0024$.

Table 4.2: Growth rates for modes $n = 6$ and $n = 7$ for PIC simulations for three values of n_b. Analytical growth rates from the solution of the dispersion relation are shown in the table captions ($\gamma_{n=6} = \gamma_{n=7} = \gamma$). Some of the simulations have been repeated.

(a) $n_b = 0.001$, $\gamma = 6.286 \cdot 10^{-3}$

np	$\gamma_{n=6}$	$\gamma_{n=7}$
8,192	$4.69{\cdot}10^{-3}$	$6.55{\cdot}10^{-3}$
8,192	$6.74{\cdot}10^{-3}$	$3.85{\cdot}10^{-3}$
16,384	$6.47{\cdot}10^{-3}$	$5.88{\cdot}10^{-3}$
32,768	$6.10{\cdot}10^{-3}$	$6.38{\cdot}10^{-3}$
32,768	$6.13{\cdot}10^{-3}$	$6.37{\cdot}10^{-3}$
32,768	$6.31{\cdot}10^{-3}$	$6.09{\cdot}10^{-3}$
65,536	$5.86{\cdot}10^{-3}$	$6.51{\cdot}10^{-3}$
131,072	$6.44{\cdot}10^{-3}$	$6.09{\cdot}10^{-3}$
262,144	$6.15{\cdot}10^{-3}$	$6.39{\cdot}10^{-3}$

(b) $n_b = 0.0024$, $\gamma = 1.479 \cdot 10^{-2}$

np	$\gamma_{n=6}$	$\gamma_{n=7}$
8,192	$1.21{\cdot}10^{-2}$	$1.51{\cdot}10^{-2}$
8,192	$1.34{\cdot}10^{-2}$	$1.49{\cdot}10^{-2}$
16,384	$1.54{\cdot}10^{-2}$	$1.39{\cdot}10^{-2}$
32,768	$1.43{\cdot}10^{-2}$	$1.50{\cdot}10^{-2}$
32,768	$1.46{\cdot}10^{-2}$	$1.49{\cdot}10^{-2}$
32,768	$1.48{\cdot}10^{-2}$	$1.47{\cdot}10^{-2}$
65,536	$1.45{\cdot}10^{-2}$	$1.50{\cdot}10^{-2}$

(c) $n_b = 0.00124$, $\gamma = 7.797 \cdot 10^{-3}$

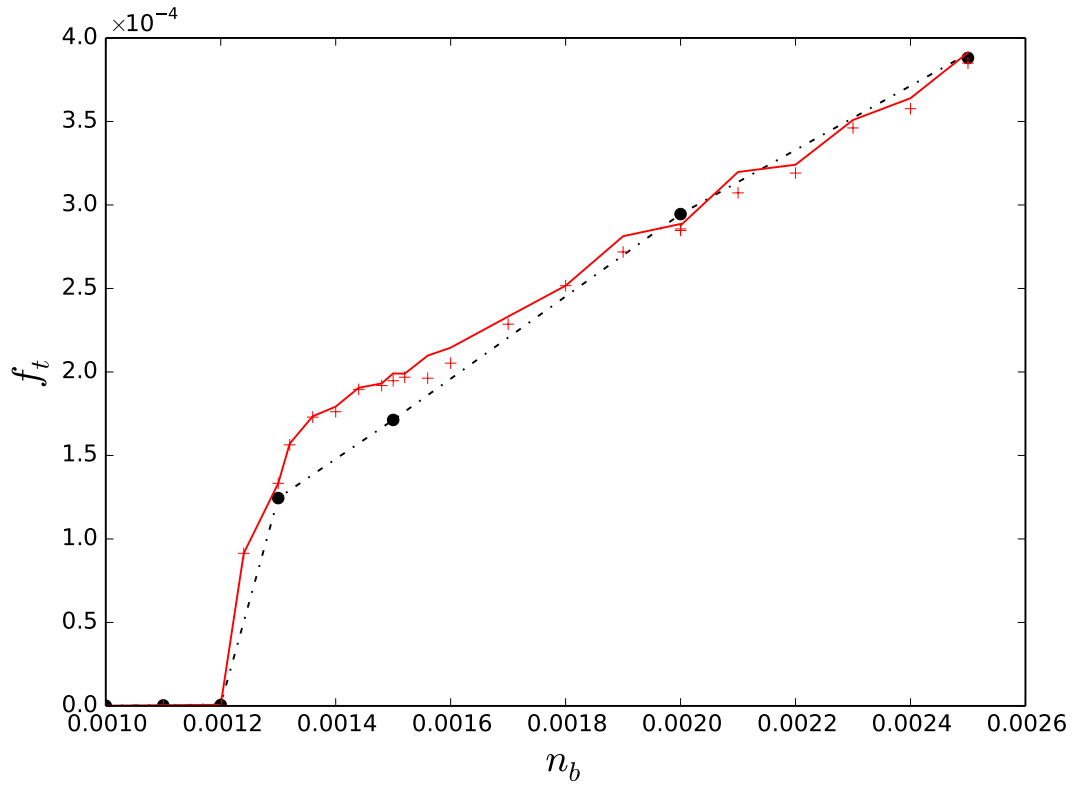| np	$\gamma_{n=6}$	$\gamma_{n=7}$
8,192	$6.68{\cdot}10^{-3}$	$7.78{\cdot}10^{-3}$
16,384	$7.16{\cdot}10^{-3}$	$7.68{\cdot}10^{-3}$
16,384	$7.89{\cdot}10^{-3}$	$7.67{\cdot}10^{-3}$
32,768	$7.82{\cdot}10^{-3}$	$7.63{\cdot}10^{-3}$
32,768	$8.27{\cdot}10^{-3}$	$7.24{\cdot}10^{-3}$
65,536	$7.33{\cdot}10^{-3}$	$8.11{\cdot}10^{-3}$
65,536	$7.71{\cdot}10^{-3}$	$7.84{\cdot}10^{-3}$
65,536	$7.85{\cdot}10^{-3}$	$7.63{\cdot}10^{-3}$
131,072	$7.69{\cdot}10^{-3}$	$7.83{\cdot}10^{-3}$
262,144	$7.66{\cdot}10^{-3}$	$7.87{\cdot}10^{-3}$
262,144	$7.77{\cdot}10^{-3}$	$7.78{\cdot}10^{-3}$
524,288	$7.88{\cdot}10^{-3}$	$7.69{\cdot}10^{-3}$

Figure 4.37: Maximum (line) and final (points) values of $f_t(t)$ as a function of the beam density n_b. Results obtained with both PIC (red) and Vlasov simulations (black).
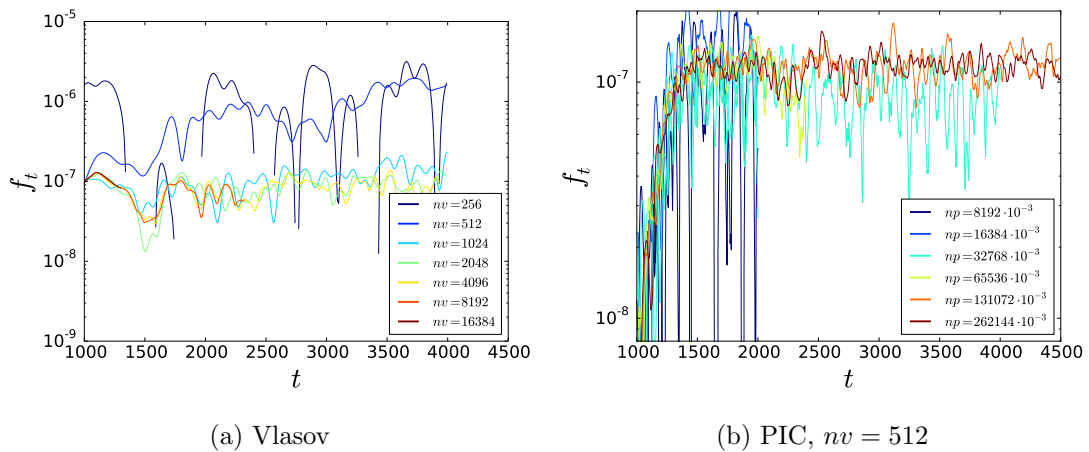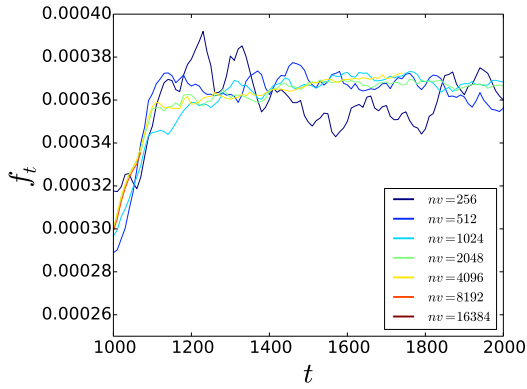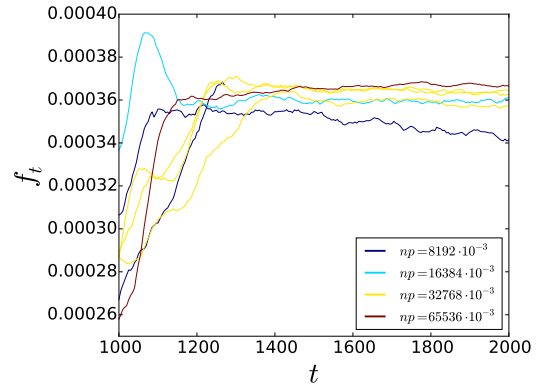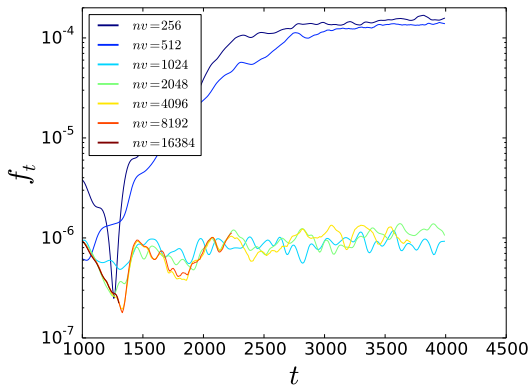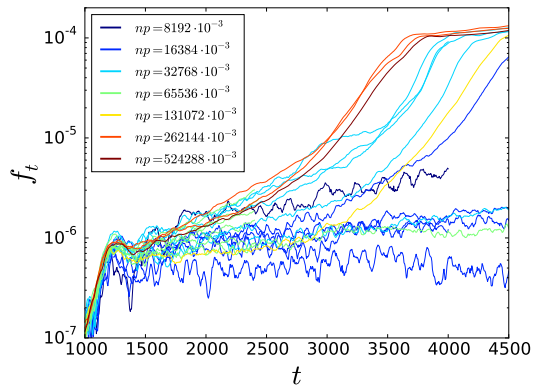

(a) Vlasov

(b) PIC, $nv = 512$

Figure 4.38: $f_t(t)$ showing convergence with increased resolution for both PIC and Vlasov simulations, for $n_b = 0.001$.
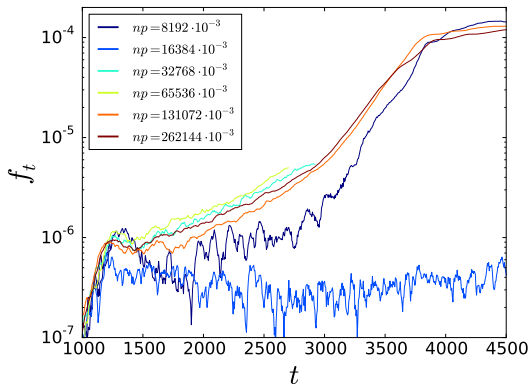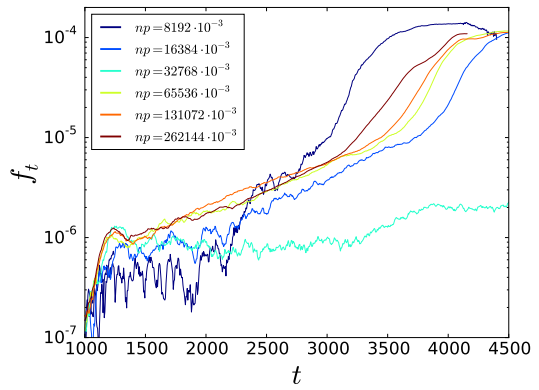
(a) Vlasov (b) PIC, $nv = 512$

Figure 4.39: $f_t(t)$ showing convergence with increased resolution for both PIC and Vlasov simulations, for $n_b = 0.0024$.

(a) Vlasov (b) PIC, $nv = 512$

(c) PIC, $nv = 1024$ (d) PIC, $nb = 2048$

Figure 4.40: $f_t(t)$ showing convergence with increased resolution for both PIC and Vlasov simulations, for $n_b = 0.00124$.

Chapter 5

Conclusions & Outlook

In §4, we have presented results from both Vlasov and PIC simulations for various kinetic systems governed by the Vlasov-Poisson set of equations.

We began by showing the linear mode decay of both strong and weak Landau damping in §4.1, for which we are able to reproduce the decay rates given by the linear dispersion relation. In §4.2, we presented results of the bump-on-tail system and we showed that the growth rates agree well with the dispersion relation. We chose two examples, one close to the maximum growth rate ($k = 0.3$), and one close to marginal stability ($k = 0.48$). We performed convergence scans for both of the example cases using both Vlasov and PIC simulations. We see that the higher growth rate case is much easier to resolve for both Vlasov and PIC simulations, and we presented an explanation of this in terms of the size of the linear mode structure. Finally, we performed a scan of the initial saturation amplitude for a range of wave numbers, and show that, at least for smaller growth rates, $A_{\text{sat}} \propto \gamma^2$. In §4.3, we began by introducing a measure of the amplitude of the n^{th} mode, doing this by showing the harmonics present in one of the previous single-mode simulations. We find good agreement with previous results, with the harmonic amplitudes following a power law throughout the linear phase. We then proceeded with simulations (again, Vlasov and PIC), of high n modes, with the number of consecutive mode numbers present (always centred around $n_0 = 20$), $n_n = \{1, 3, 5\}$. We measured the growth rates of the modes, and presented convergence plots for the growth rates as a function of the resolution. We claim that the number of modes present in the simulation has no significant effect on the resolution requirements of the simulation.

Finally, in §4.4, we presented a number of results for a two-mode system spanning the threshold at which the islands (as measured in single-mode cases) begin to overlap. We controlled the system by varying the beam density, whilst keeping the growth rates of the two modes fixed to each other by fine-tuning the simulation domain size ($k_n = nk_0$, $k_0 = 2\pi/L$). We show how a sample of the particle trajectories from PIC simulations can be used to form a useful diagnostic in terms of following the islands' width and structure. We see that when the growth rate is sufficiently small and the island widths are sufficiently small in velocity space such that they do not overlap, then the two-mode system behaviour is almost exactly that of the one mode systems and the modes evolve independently of each other, saturating at the same amplitudes and at the same time as in the single-mode systems. However, as the growth rates, and therefore also the island widths, were increased such that the islands overlap (as measured in single-mode cases), we see a qualitatively different behaviour. We see that the islands combined and the regions of phase space became mixed, and the distribution function $f(v)$ flattened across the entire overlap region. We see a significant increase in the final

mode amplitudes compared to the single-mode systems, and we demonstrated the insensitivity of this effect to initial mode amplitudes. We quantified this redistribution effect by looking at the total amount of the phase space density which was transported across a particular value of velocity located between the islands. We showed the dependence of this transported phase space density on the beam density for both Vlasov and PIC simulations and find matching results showing a discontinuity at the threshold over overlap. We also presented resolution scans at each of three sample cases (one non-overlapping, one strongly overlapping, and one marginally overlapping). We find that in the cases of non-overlapping and strongly overlapping, the resolution requirements for both Vlasov and PIC simulations were relatively modest. However, for the marginally overlapping case, the requirements were more stringent.

We also notice the interesting observation that, in the marginally overlapping case, Vlasov simulations with insufficient velocity grid resolution yielded values of the transport metric two orders of magnitude larger than those with increased resolution (this agreed with predictions from cases which do and do not exhibit overlap respectively). Therefore we conclude that, with insufficient resolution, the Vlasov simulation is unable to resolve small separation between phase space structures, and these structures are spuriously connected. PIC simulations of the same case also yielded these same two different values of the transport metric spanning two orders of magnitude. However, for PIC simulations, we found that an increase in the number of particles increases the probability of yielding the higher (corresponding to overlapping) case, whereas simulations with smaller number of particles were more likely to yield the lower, non-overlapping value (in constrast with the Vlasov observation). We surmise that this behaviour in the PIC simulations is related to the probabilistic nature of the PIC method, with only a finite number of the physical particle trajectories are present. Therefore, if an overlap is physically present, but small, then for the overlap to be experienced by the simulation, some of the critical trajectories must be present. Therefore, as the number of markers in the simulation is increased, the probability of sampling these critical trajectories is also increased.

In summary, in this work, we have presented simulation results of the one-dimensional bump-on-tail system using both of the leading types of numerical simulation method. We have focussed on the case of the multi-mode system, and have shown the change in behaviour seen when two modes cross the threshold of overlap. We have also focussed on the question of numerical requirements and resolution, for which the one-dimensional model is a very useful tool due to its relatively inexpensive simulation nature (two-dimensional phase space). We have tried to draw some conclusions of numerical requirements and their link to physical scales, which we hope may be a helpful starting point for those performing similar multi-mode simulations with higher-dimensional systems (e.g. five-dimensional phase space in gyrokinetic simulations).

Outlook

Following on from the work presented here, we would propose the following courses of future action.

As we have observed phase space transport for overlapping multi-mode systems, a natural question relates to the mechanisms of the transport. An interesting result would be to quantify the diffusion coefficient across phase space, and to see how this value depends on phase space coordinates and on system parameters. The diffusion coefficients from similar overlapping mode systems in the three-dimensional realistic geometry of HAGIS simulations could then also be calculated, and the coefficients compared. We

would then be interested to see how similar the behaviour is in the two cases.

With future access to a so-called 1D2V code, in which there exists one dimension in real space and two dimensions of velocity space, similar simulations could be performed, but with the added physics that only a fraction of phase space with a particular velocity would be resonant with the mode. This would allow a study of the role of resonant and non-resonant particles.

With the inclusion of more modes and modes of higher mode numbers, which would also likely be the case in ITER plasmas, a move towards comparison with more realistic ITER scenarios can be performed. This can be combined with the calculation of diffusion coefficients, where we could explore the differences between a system of many nearby marginally unstable resonances with the case of the dual resonance, or the highly unstable case.

We made the conclusion that PIC methods are liable to underpredict transport in marginal cases in the case of an insufficient number of markers, if critical trajectories are not sampled. We would like to test this hypothesis by creating a synthetic test case, and systematically constructing various phase space samplings in an attempt to demonstrate this effect. Depending on the result of this test, we would like to try to reproduce the effect with the marker loading in HAGIS, and to investigate whether a more efficient loading of markers can be found with certain prior knowledge of a system. It should also be investigated whether live diagnostics can be created to monitor the distribution of the markers.

Acknowledgements

Bibliography

[1] Tables of physical & chemical constants (16th edition 1995). 4.7.4 Nuclear fusion. Kaye & Laby online. Version 1.0, 2005.

[2] J. A. Wesson. *Tokamaks.* Oxford University Press, 3rd edition, 2004.

[3] A. B. Mikhailovskii. Thermonuclear "drift" instabilities. *Journal of Experimental and Theoretical Physics*, 1975.

[4] M. N. Rosenbluth and P. H. Rutherford. Excitation of alfvén waves by high-energy ions in a tokamak. *Phys. Rev. Lett.*, 34:1428–1431, Jun 1975.

[5] H. Alfvén. Existence of electromagnetic-hydrodynamic waves. 150:405–406, October 1942.

[6] L. Chen. Alfvn waves: a journey between space and fusion plasmas. *Plasma Physics and Controlled Fusion*, 50(12):124001, 2008.

[7] D. J. Sigmar, C. T. Hsu, R. White, and C. Z. Cheng. Alpha-particle losses from toroidicity-induced Alfvn eigenmodes. part ii: Monte Carlo simulations and anomalous alpha-loss processes. *Physics of Fluids B*, 4(6):1506–1516, 1992.

[8] S. E. Sharapov, B. Alper, H. L. Berk, D. N. Borba, B. N. Breizman, C. D. Challis, I. G. J. Classen, E. M. Edlund, J. Eriksson, A. Fasoli, E. D. Fredrickson, G. Y. Fu, M. García-Mu noz, T. Gassner, K. Ghantous, V. Goloborodko, N. N. Gorelenkov, M. P. Gryaznevich, S. Hacquin, W. W. Heidbrink, C. Hellesen, V. G. Kiptily, G. J. Kramer, Ph. Lauber, M. K. Lilley, M. Lisak, F. Nabais, R. Nazikian, R. Nyqvist, M. Osakabe, C. Perez von Thun, S. D. Pinches, M. Podesta, M. Porkolab, K. Shinohara, K. Schoepf, Y. Todo, K. Toi, M. A. Van Zeeland, I. Voitsekhovich, R. B. White, V. Yavorskij, ITPA EP TG, and JET-EFDA Contributors. Energetic particle instabilities in fusion plasmas. *Nuclear Fusion*, 53(10):104022, 2013.

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

[10] M. García-Muñoz, N. Hicks, R. van Voornveld, I. G. J. Classen, R. Bilato, V. Bobkov, M. Brüdgam, H.-U. Fahrbach, V. Igochine, S. Jaemsae, M. Maraschek, and K. Sassenberg. Convective and diffusive energetic particle losses induced by shear Alfvén waves in the ASDEX Upgrade tokamak. *Phys. Rev. Lett.*, 104:185002, May 2010.

[11] S. D. Pinches, I. T. Chapman, Ph. Lauber, H. J. C. Oliver, S. E. Sharapov, K. Shinohara, and K. Tani. Energetic ions in ITER plasmas. *Physics of Plasmas (1994-present)*, 22(2):–, 2015.

[12] Ph. Lauber. Local and global kinetic stability analysis of Alfvén eigenmodes in the 15 MA ITER scenario. *Plasma Physics and Controlled Fusion*, 57(5):054011, 2015.

[13] Ph. Lauber. Super-thermal particles in hot plasmas – kinetic models, numerical solution strategies, and comparison to tokamak experiments. *Physics Reports*, 533:33–68, December 2013.

[14] H. L. Berk and B. N. Breizman. Saturation of a single mode driven by an energetic injected beam. III. Alfvén wave problem. *Physics of Fluids B*, 2:2246–2252, September 1990.

[15] H. L. Berk, B. N. Breizman, and H. Ye. Scenarios for the nonlinear evolution of alpha-particle-induced Alfvén wave instability. *Phys. Rev. Lett.*, 68:3563–3566, June 1992.

[16] H. L. Berk, B. N. Breizman, J. Fitzpatrick, and H. V. Wong. Line broadened quasi-linear burst model. *Nuclear Fusion*, 35(12):1661, 1995.

[17] R. G. L. Vann, R. O. Dendy, G. Rowlands, T. D. Arber, and N. d'Ambrumenil. Fully nonlinear phenomenology of the Berk-Breizman augmentation of the Vlasov-Maxwell system. *Physics of Plasmas*, 10(3):623–630, 2003.

[18] M. K. Lilley, B. N. Breizman, and S. E. Sharapov. Effect of dynamical friction on nonlinear energetic particle modes. *Physics of Plasmas*, 17(9):–, 2010.

[19] M. Lesur. *The Berk-Breizman Model as a Paradigm for Energetic Particle-driven Alfvén Eigenmodes*. PhD thesis, Ecole Doctorale de l'Ecole Polytechnique, 2011.

[20] M. Schneller, Ph. Lauber, M. Brüdgam, S. D. Pinches, and S. Günter. Double-resonant fast particle-wave interaction. *Nuclear Fusion*, 52(10):103019, 2012.

[21] A. A. Vlasov. On vibration properties of electron gas. *Journal of Experimental and Theoretical Physics*, 3:291, 1938.

[22] L. D. Landau. On the vibration of the electron plasma. *J. Phys. (U.S.S.R.)*, 10, 1946.

[23] T.H. Stix. *The Theory of Plasma Waves*. McGraw-Hill, 1962.

[24] M. Brambilla. *Kinetic Theory of Plasma Waves: Homogeneous Plasmas*. International series of monographs on physics. Clarendon Press, 1998.

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

[26] R. Z. Sagdeev and A. A. Galeev. *Nonlinear Plasma Theory*. 1969.

[27] M. Kraus. *Variational Integrators in Plasma Physics*. PhD thesis, Technische Universität München, 4 2013.

[28] M. Kraus and O. Maj. Variational integrators for nonvariational partial differential equations. *Physica D*, 310:37 – 71, 2015.

[29] M. Kraus, O. Maj, and E. Sonnendrücker. Variational integrators for the Vlasov-Poisson system. In Preparation.

[30] J. E. Marsden, G. W. Patrick, and S. Shkoller. Multisymplectic geometry, variational integrators, and nonlinear PDEs. *Communications in Mathematical Physics*, 199:351 – 395, 1998.

[31] J. E. Marsden and M. West. Discrete mechanics and variational integrators. *Acta Numerica*, 10:357–514, 2001.

[32] C. Z. Cheng and G. Knorr. The integration of the Vlasov equation in configuration space. *Journal of Computational Physics*, 22:330–351, November 1976.

[33] F. Filbet and E. Sonnendrücker. Comparison of Eulerian Vlasov solvers. *Computer Physics Communications*, 150(3):247 – 266, 2003.

[34] F. H. Harlow. A machine calculation method for hydrodynamic problems. Technical report, Los Alamos Scientific Laboratory, 1955.

[35] W. Deng and G. Y. Fu. Optimization by marker removal for particle simulations. *Computer Physics Communications*, 185(1):96 – 105, 2014.

[36] W. Deng. PIC1D-PETSc, 2012. [Online; accessed August 2015].

[37] T. M. O'Neil and J. H. Malmberg. Transition of the dispersion roots from beam-type to Landau-type solutions. *Physics of Fluids*, 11:1754–1760, August 1968.

[38] T. Tajima and F. W. Perkins. Sherwood Theory Meeting, 1983.

[39] A. M. Dimits and W. W. Lee. Partially linearized algorithms in gyrokinetic particle simulation. *J. Comput. Phys.*, 107(2):309–323, August 1993.

[40] W. Deng. 1D electrostatic δf particle-in-cell (PIC) formulation in vector-matrix form. Available online, May 2014.

[41] T. M. O'Neil, J. H. Winfrey, and J. H. Malmberg. Nonlinear interaction of a small cold beam and a plasma. *Physics of Fluids*, 14(6):1204–1212, 1971.