EUROFUSION WPJET1-PR(15) 14358

J Leddy et al.

On the validity of drift-reduced fluid models for tokamak plasma simulation

Preprint of Paper to be submitted for publication in
Plasma Physics and Controlled Fusion

This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014-2018 under grant agreement No 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission.

On the validity of drift-reduced fluid models for tokamak plasma simulation

Jarrod Leddy[1,2], Ben Dudson[1], Michele Romanelli[2], and JET Contributors[3]
E-mail: jbl504@york.ac.uk
[1]York Plasma Institute, Department of Physics, University of York, Heslington YO10 5DD, UK
[2]CCFE, Culham Science Centre, Abingdon, Oxfordshire OX14 3DB, UK
[3]EUROfusion Consortium, JET, Culham Science Centre, Abingdon, OX14 3DB, UK

Abstract

Drift-reduced plasma fluid models are commonly used in plasma physics for analytic studies and simulations, so the validity of such models must be verified for the regions of parameter space in which tokamak plasmas exist. By deriving and comparing the linear dispersion relations for the drift-wave instability for both a drift-reduced model and a full-velocity model, the importance of the physics lost with the drift-reduction is examined. This analysis is generalised for typical tokamak parameter spaces and is then applied directly to JET data. It is found that drift-reduced models are generally more applicable to the edge plasma ($< 10\%$ error), while the core plasma shows more significant disagreement ($> 30\%$ error) particularly at mid-radius. The effect of drift-wave mode number and wavelength also play a key role in determining the accuracy of drift-reduced models.

1 Introduction

Fluid models are often used to describe plasma behaviour in a magnetic field, especially utilising the closure developed by Braginskii [1] that is valid for highly collisional plasmas such as linear devices and in the tokamak edge where collisional damping is the dominate damping mechanism. It is often asserted that only kinetics and gyrokinetics can truly describe plasma dynamics in collisionless regimes, such as the core of tokamaks [2]. This is because collisional damping plays a strong role in the formation of turbulence, but in its absence in collisionless plasmas, finite Larmor radius and kinetic effects such as Landau damping dominate instead. These are analytically present only in gyrokinetic [3] and sometimes approximated in gyro-Landau fluid models [4, 5]; however, core turbulent transport has been qualitatively and quantitatively reproduced using a pure fluid description as well [6]. Due to the complicated and expensive nature of gyrokinetic models, it is preferable, when appropriate, to utilise fluid models instead. The aim of this work is not to test the regimes in which fluid models are applicable, but to instead review the validity of drift-reduced fluid models in the regimes where the fluid approximation is reasonable.

There has been a large effort to derive fluid models that provide both corrections and simplifications to the original Braginskii system [7, 8]. One such simplification, the so-called drift-reduction, was derived by Mikhailovskii and Tsypin [7] and is a slow ordering that assumes $\omega < \omega_{ci}$ and $\rho_i = 0$. The drift-reduction is a method of simplifying the momentum equation by taking its curl, resulting in an equation for the evolution of vorticity, $\vec{W} = \vec{\nabla} \times \vec{v}$. For this new system to be closed, an assumption is made that the perpendicular velocities are dominated by the $\vec{E} \times \vec{B}$ drift, which relates the parallel vorticity to the potential: $W_{\parallel} = \nabla_{\perp}^2 \phi$ (cgs Gaussian units are used for the duration of the paper).

The drift-reduction requires that perpendicular force balance be satisfied at all times, so perpendicular acceleration terms are neglected and fast waves, such as the fast magnetosonic wave, are removed from the system. The effect of the pressure gradient on the velocity evolution is not carried into the vorticity equation as it disappears with the curl. This is especially significant because the largest pressure gradients will be perpendicular to the magnetic field, and it is the perpendicular velocity equations that are replaced by an equation for the vorticity. This approximate treatment of the pressure dynamics leads to many of the differences between the full-velocity and drift-reduced systems.

By using linearisation techniques, the behaviours of these models are compared in order to determine in which cases the drift-reduction is acceptable. This is feasible because tokamaks operate in a well-defined yet broad parameter space that can be explicitly explored. Both the reduced and full systems are simplified to the incompressible limit (such that $\vec{\nabla} \cdot \vec{v} = 0$) to look at the most basic case that still produces

drift-waves, which are the most universal drive mechanism for tokamak plasma turbulence since they require only a pressure or density gradient and finite resistivity to exist [9, 10]. Any differences for this simplified case are thus fundamental and will carry on into more complex scenarios. Though turbulence is a thoroughly non-linear phenomenon, the linear growth rate of the drive instability indicates stability and determines the non-linear saturation timescale and quasilinear flux making linear analysis of these instabilities both relevant and essential.

2 Full velocity vs drift-reduced models

A full-velocity model is one that evolves all three components of equation 1, the ion momentum equation [11],

$$mn\left(\frac{\partial\vec{v}}{\partial t}+\vec{W}\times\vec{v}\right)=\frac{\vec{J}\times\vec{B}}{c}-\nabla p - \frac{mn}{2}\nabla(\vec{v}\cdot\vec{v})-mn\chi_v(\vec{\nabla}\times\vec{W}) \tag{1}$$

where m is the mass of the ions, n is the ion density, $\vec{v}$ is the ion velocity, $\vec{W}$ is the ion vorticity, $\vec{J}$ is the current density, $\vec{B}$ is the magnetic field, p is the total pressure, χ_v is the velocity diffusivity, and c is the speed of light. By taking the curl of equation 1, an equation for vorticity is obtained. It is convenient to take the parallel component of the vorticity equation, as shown in equation 2, because it includes the behaviour of the perpendicular velocities:

$$W_\parallel=\hat{b}\cdot\left(\vec{\nabla}\times\vec{v}\right)=\left(\frac{\partial v_x}{\partial z}-\frac{\partial v_z}{\partial x}\right)\hat{y} \tag{2}$$

where $\hat{x}$ and $\hat{z}$ are the perpendicular directions and $\hat{y}$ is parallel to the magnetic field line. This geometry will be explained in more detail in the next section. The normalised drift reduced equations, as derived by Hazeltine, et al. [12], can then be written in the incompressible limit as follows:

$$\begin{aligned}\frac{\partial p}{\partial t}&=-\left[\phi,p\right]\\ \frac{\partial W_\parallel}{\partial t}&=-\left[\phi,W_\parallel\right]-\nabla_\parallel J_\parallel\\ \frac{\partial A_\parallel}{\partial t}&=-\nabla_\parallel\phi+\eta J_\parallel+\nabla_\parallel p\end{aligned} \tag{3}$$

where p is the pressure, ϕ is the electric potential, $A_\parallel$ is the parallel vector potential, $[f,g]=\frac{\partial f}{\partial x}\frac{\partial g}{\partial z}-\frac{\partial f}{\partial z}\frac{\partial g}{\partial x}$ are the standard advection brackets, $J_\parallel=\nabla_\perp^2 A_\parallel$, $W_\parallel=\nabla_\perp^2\phi$, and η is the parallel resistivity. The equation for $v_\parallel$ does not couple to these in the incompressible limit so is omitted, but it is important to note that the parallel velocity will evolve to maintain $\vec{\nabla}\cdot\vec{v}=0$.

Using the same normalisations, the full-velocity model in the incompressible limit is given by

$$\begin{aligned}\frac{\partial p}{\partial t}&=\nabla p\cdot\vec{v}\\ \frac{\partial v_x}{\partial t}&=\left(\vec{J}\times\vec{B}\right)_x-\nabla_x p\\ \frac{\partial v_z}{\partial t}&=\left(\vec{J}\times\vec{B}\right)_z-\nabla_z p\\ \frac{\partial A_\parallel}{\partial t}&=-\nabla_\parallel\phi+\eta J_\parallel+\nabla_\parallel p\\ \frac{\partial A_z}{\partial t}&=0\;\;=\;\;-\nabla_z\phi+\eta_\perp J_z+\nabla_z p\\ &\quad+\left(\vec{v}\times\vec{B}\right)_z+\left(\vec{J}\times\vec{B}\right)_z\end{aligned} \tag{4}$$

with $\vec{\nabla}\cdot\vec{J}=0$, $\vec{J}=\vec{\nabla}\times\vec{B}$, and $\vec{B}=\vec{\nabla}\times\vec{A}$. To evolve the vector potential, the generalised Ohm's law given by Lifshitz [13] is used with temperature gradients neglected due to the isothermal assumptions. The equation for A_x is excluded above because it simply evolves to maintain force balance without coupling to the remaining equations. For both systems parallel derivatives are taken along the perturbed field by defining $\nabla_\parallel f=\partial_\parallel f-\left[A_\parallel,f\right]$. It is important when exploring the effects of the drift-reduction to ensure the two systems (full-velocity and drift-reduced) are identical in all other aspects. To do this, the full-velocity system was drift-reduced and in the linear limit reproduces exactly the dispersion relation of the Hazeltine model, which will be shown in the next section in equation 5. In this way, the effects of the gyro-viscous cancellation which is used in both models are not observed in our comparison.

3 Linearisation

For all linearisations, we define a quasi-3D, orthogonal coordinate system (x-y-z) such that the equilibrium magnetic field B_0 is in the y-direction, the equilibrium current density J_0 is in the negative z-direction, and the background pressure gradient is in the x-direction; however, perturbations are only

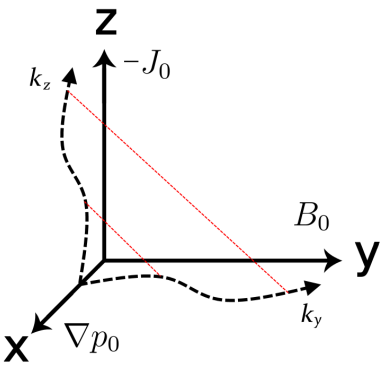

Figure 1: The geometry for the linearisation is quasi-3D with equilibrium pressure gradient, current density, and magnetic field that satisfy force balance. Perturbations are in y and z such that the total perturbation is at an angle to the magnetic field, B_0.

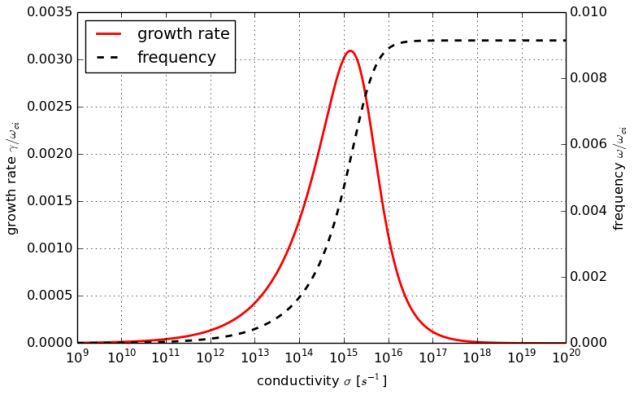

Figure 2: Growth rate (dashed) and frequency (solid) of the drift-wave instability as a function of conductivity for the drift-reduced system. Frequency, growth rate, and conductivity are all normalised to the ion cyclotron frequency.

in y and z with no extent in x consistent with a local approach, as detailed in figure 1.

The background pressure gradient is present to drive the drift-wave instability, and the background current density and magnetic field are provided to satisfy force balance. All perturbations are of the form $\tilde{f} = e^{i(k_y y + k_z z - \Omega t)}$ where Ω is the complex frequency defined as $\Omega = \omega + i\gamma$. Both systems were converted to cgs units prior to linearisation so that physical parameters, such as Alvén speed, could be more easily substituted into the resulting dispersion relations.

3.1 Drift-reduced dispersion relation

The drift-reduced system in equation 3, once linearised, results in the following dispersion relation:

$$\Omega^3 + \left(\omega_* + i\eta\frac{v_A^2 k_z^2 \omega_{pi}^2}{4\pi\omega_{ci}^2}\right)\Omega^2 - \left(v_A^2 k_y^2\right)\Omega - \left(v_A^2 k_y^2 \omega_*\right) = 0 \tag{5}$$

where ω_* is the drift-wave frequency, ω_{ci} is the ion cyclotron frequency, ω_{pi} is the ion plasma frequency, and v_A is the Alfvén speed. These are defined as

$$\omega_* = \frac{|\nabla p_0| k_z}{m_i n_0 \omega_{ci}} \qquad \omega_{ci} = \frac{eB_0}{m_i c}$$

$$\omega_{pi} = \sqrt{\frac{4\pi n_0 e^2}{m_i}} \qquad v_A = \frac{B_0^2}{\sqrt{4\pi m_i n_0}}.$$

Parallel Alfvén waves as well as resistive drift-waves can be seen in the terms of the dispersion relation. In the case of zero resistivity the waves are stable and simply propagate. For $\eta > 0$ the most unstable growth rate and corresponding frequency can be extracted using typical values for magnetic field, pressure gradient, and background density within a tokamak ($B = 10^4$G, $|\nabla p| = 10^6$Ba/cm, $n_0 = 10^{12}\text{cm}^{-3}$, $T = 50$eV, $k_y = 0.03\text{cm}^{-1}$, and $k_z = 10\text{cm}^{-1}$).

These values are also chosen to satisfy $\omega_* < \omega_{ci}$ such that the ion cyclotron frequency is the highest frequency in the system (note that the pressure gradient here is typical for the pedestal and will be lower in other areas of the tokamak, reinforcing this ordering). The resulting frequency and growth rate are plotted in figure 2 as a function of conductivity, $\sigma = 1/\eta$ (ie. the inverse of resistivity) with typical values of 10^{13}s^{-1} in the pedestal.

3.2 Full-velocity dispersion relation

When the full-velocity system is linearised, the resulting dispersion relation (equation 6) contains the extra, fast-physics that was lost in the drift-reduction.

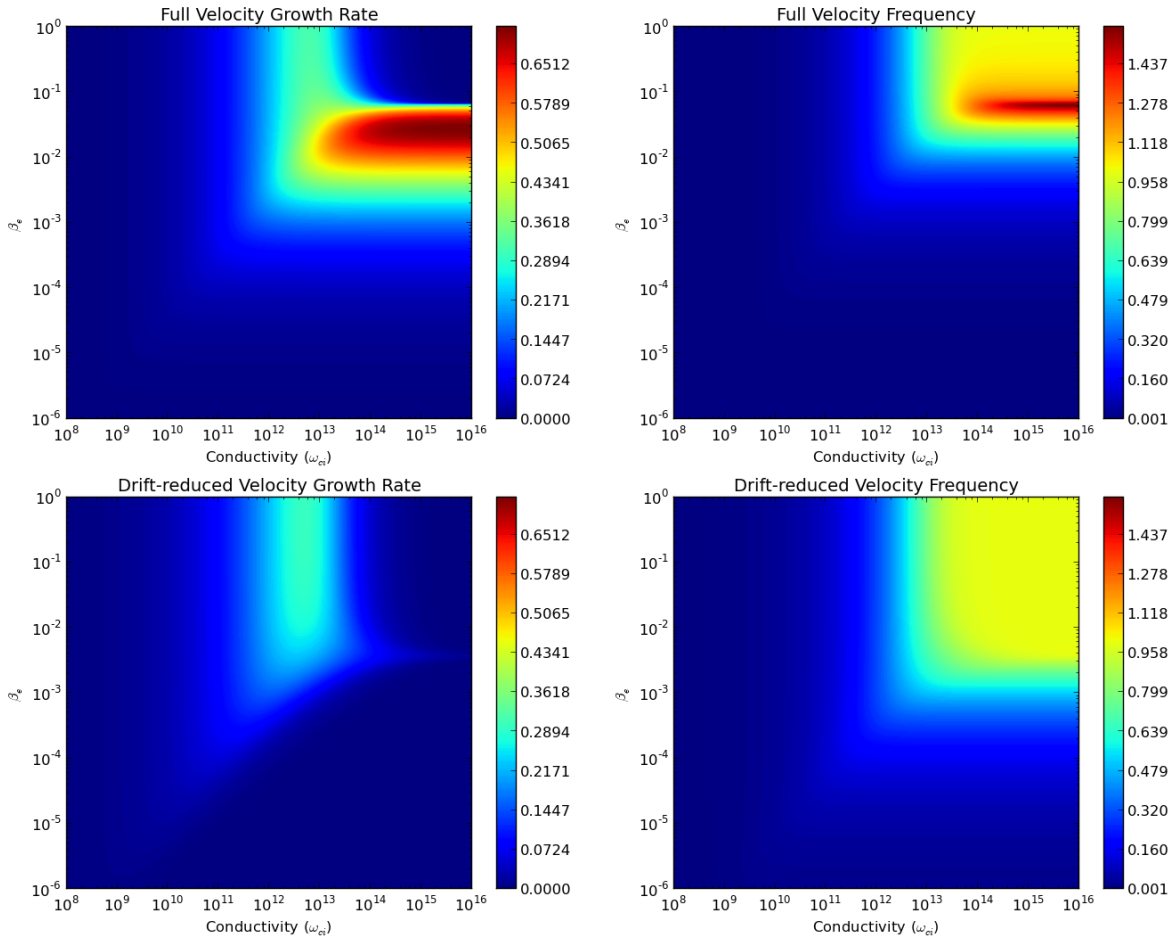

Figure 3: Full velocity and drift-reduced growth rates and frequencies as a function of conductivity and β_e at $B_0 = 1\text{T}$. Growth rates and frequencies are normalised to the drift-wave frequency, ω_*.

$$\left(\frac{\omega_{pi}^2\eta}{2\pi i\omega_{ci}^2}\right)\Omega^4+\left(1+\frac{v_A^2k_y^2}{\omega_{ci}^2}+\frac{\omega_{pi}^4v_A^2\left(k_y^2+2k_z^2\right)\eta^2}{16\pi^2\omega_{ci}^4}+\frac{\omega_{pi}^2\omega_*\eta}{\pi i\omega_{ci}^2}\right)\Omega^3 + \left(\omega_*+i\eta\frac{v_A^2\left(4k_y^2+k_z^2\right)\omega_{pi}^2}{4\pi\omega_{ci}^2}\right)\Omega^2-\left(v_A^2k_y^2\right)\Omega-\left(v_A^2k_y^2\omega_*\right)=0 \tag{6}$$

The perpendicular resistivity has been approximated to be double the parallel resistivity, as given in Wesson [14]. Notice this expression has a higher order in Ω compared with equation 5 due to the additional equation for the perpendicular velocities, which results in an extra mode in the growth rate and frequency. The two dispersion relations, equations 5 and 6, are identical when Ω^4 term, the last three terms in Ω^3, and the parallel wave number in the second term of the Ω^2 are neglected, indicating that these terms contain the physics lost in the drift-reduction. This includes various propagating parallel and perpendicular resistive modes and the ion cyclotron wave.

For the same values of magnetic field, background density, and pressure gradient the growth rate and frequency of this expression are quite similar to that of the drift-reduced model, so it was necessary to look at the solutions over a large parameter space in conductivity and electron beta, defined by

$$\beta_e=\frac{p_{gas}}{p_{mag}}=\frac{8\pi nT_e}{B^2}.$$

In figures 3 and 4 the magnetic field is set to constant $B = 1$T and the density is adjusted to vary beta. This is useful to do because the terms in equation 6 are not functions of only β_e - they depend on various combinations of density and magnetic field. In essence, the parameter space is 3D, however this is not easily visualised so magnetic field has been held constant for illustrative purposes. An important difference between the qualitative behaviour of the two models is demonstrated by the monotonicity of the full velocity growth rate as a function of conductivity for $\beta_e < 10^{-2}$, while the drift-reduced growth rate is always peaked (ie. stabilised at high conductivity). This is due to the fast magnetosonic wave resonating with and destabilising the drift-waves, even at low resistivity. The cut-off where the low frequency

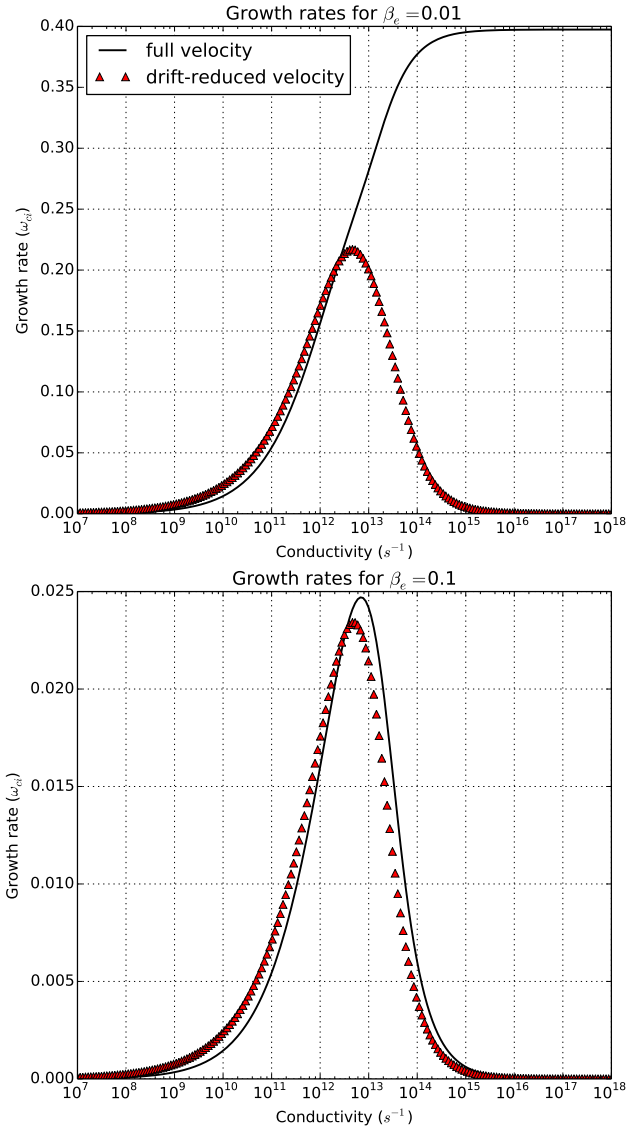

Figure 4: Full velocity and drift-reduced growth rates as a function of conductivity for two values of electron beta, $\beta_e = 0.01$ (left) and $\beta_e = 0.1$ (right).

assumption breaks down is at $\beta_e = 0.079$ for these parameters, at which $\omega_* \sim \omega_{ci}$ and below which the drift-reduced is expected be inaccurate.

4 Tokamak relevance

The parameter space in which tokamaks operate is specific to the region within the tokamak (core vs edge) and the particular tokamak in question. For a large tokamak of size similar to JET, the Joint European Torus at the Culham Science Centre, the core operates around $\beta_e = 0.03$ and $\sigma = 10^{15}\text{s}^{-1}$, while in the edge $\beta_e = 0.005$ and $\sigma = 10^{12}\text{s}^{-1}$.

Figure 5 depicts the percent difference in growth rates between the full-velocity and drift-reduced systems as given by

$$\Delta_{\%} = \left| \frac{\gamma_{FV} - \gamma_{DR}}{\gamma_{FV}} \right|. \tag{7}$$

It is clear that at low conductivity the drift-reduction breaks down for all values of β_e. This is due to the terms exclusively in equation 6 that are functions of η and η^2 becoming very large at low conductivity, $\sigma = \eta^{-1}$.

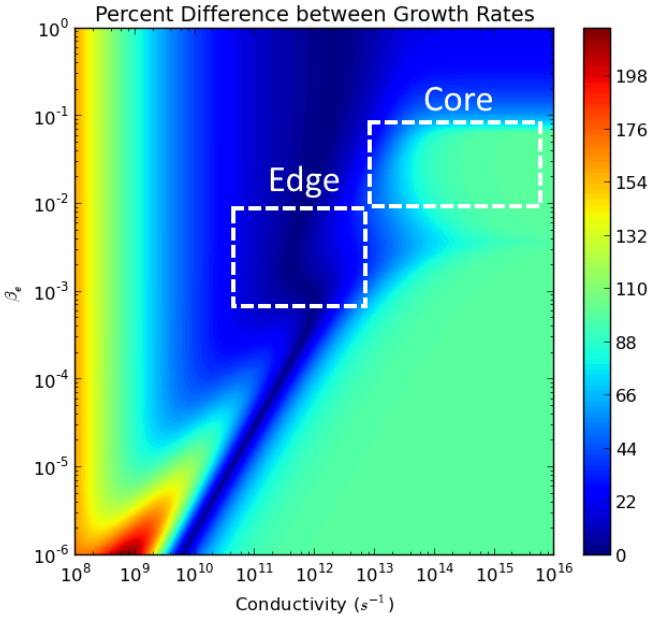

Figure 5: The percent difference from equation 7 as a function of conductivity and electron beta. Usual operational regimes for the core and edge in JET are marked.

At low β_e, which corresponds to low density, we see a fairly universal disagreement between the models. Since $\omega_* \propto n^{-1}$ and $v_A \propto n^{-1/2}$ these plasma parameters become larger at low density. The plasma frequency $\omega_{pi} \propto n^{1/2}$, so it becomes small at low density, while the ion cyclotron frequency is not a function of density. All of the additional terms in equation 6 vanish at low density except for the second term of the Ω^3 term, $\frac{v_A^2 k_y^2}{\omega_{ci}^2}$ which is proportional to n^{-1}. At low density and high conductivity this term dominates, but as conductivity is lowered, the η^2 term takes over, thus the small area of agreement even at low β_e.

The indication is that drift-reduced models are able to accurately reproduce edge behaviour, where errors can be as low as 0%. There are regions even in the edge, however, where a drift-reduced model may not be appropriate and errors can reach as high as 100%. In the core, there is a fairly consistent error of around 100% from the full-velocity solution. Such analysis, however, is based on parameters chosen for the example case shown in figure 5, such as a constant pressure gradient $\nabla p = 10^6$ Ba/cm and temperature $T_e = 50eV$. The parameter space is multi-dimensional, so to reduce the number of free parameters as much as possible, experimental data is examined.

4.1 Analysis with JET data

Using the high resolution Thomson scattering (HRTS) system on JET, density and temperature radial profiles have been acquired for shot 87045. The time trace of the density and temperature in the core are shown in the top plot of figure 6, which indicates two regimes of interest: L-mode and H-mode (black and green, respectively, vertical dotted lines). Some parameters are then calculated from these profiles, such as the parallel resistivity (assumed Spitzer [11])

$$\eta_S = \frac{2\sqrt{2\pi m_e} e^2 \ln \Lambda}{3T_e^{3/2}} \tag{8}$$

with the Coulomb logarithm approximated as $\ln \Lambda \simeq 14.9 - 0.5 \ln n_e + \ln T_e$ [14]. The only assumed values are the parallel and perpendicular wave numbers $k_y = 0.03\text{cm}^{-1}$ and $k_z = 10.0\text{cm}^{-1}$, which have been chosen based on experimental and theoretical evidence [16]. The dependence of the results on this choice is examined in section 4.2. By solving the dispersion relations for each system the percentage error between the two models is compared in the bottom right plot of figure 6.

A slightly different behaviour is seen for the actual JET data compared to the more general results,

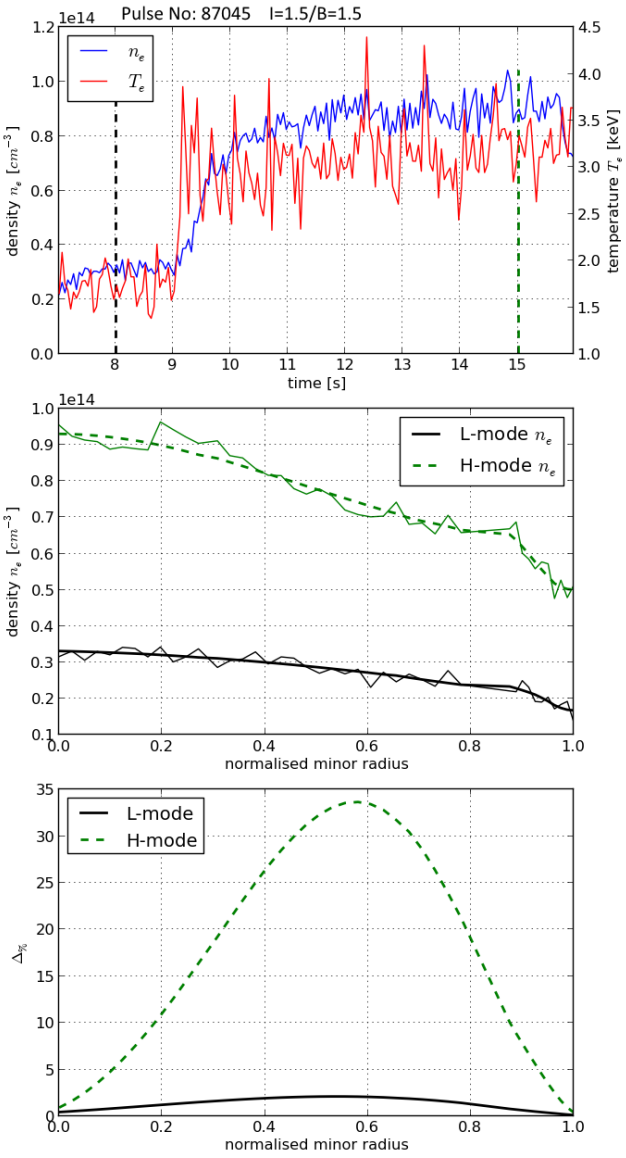

Figure 6: HRTS data from JET (pulse number 87045) is used to calculate how well the drift-reduced model describes the linear drift-wave growth rate. A time trace (top) of the core density and temperature is shown with vertical lines marking the L-mode (black) and H-mode (green) that are investigated individually. Density profiles show increase in confinement (ie. core density) and development of the pedestal (bottom left). The error between drift-reduced and full-velocity models is shown (bottom right) to be a function of the radial position within the plasma, where the far edge and deep core are shown to be most accurate. Toroidal flux values from EFIT [15] are used to estimate the normalised minor radius.

mostly due to the pressure gradient having a profile instead of being held constant. In the deep core, the pressure gradient approaches zero as the pressure reaches a maximum, reducing the drift-wave drive to nearly zero and stabilising both models. This causes accuracy to be recovered in the core, where previously it was seen to be very inaccurate, however the inaccuracies persist in the outer core peaking at $r_N \approx 0.58$cm. The edge remains accurate, especially approaching the scrape-off-layer where the plasma is highly collisional. The HRTS radial resolution is not high enough to describe the pedestal in detail, but even at the top of the pedestal ($r_N \approx 0.9$) the error is seen to be approximately 10%. The increase in error of the H-mode from 0% up to $\sim 10\%$ going into the core is due to the inflection of the pressure and density profile fits, which are likely un-physical characteristics.

The L-mode analysis shows much lower disagreement between the drift-reduced and full-velocity systems than the H-mode, remaining under 5% across the entire plasma profile. This corresponds to the drift-wave frequency remaining an order of magnitude lower than the cyclotron frequency, a condition of the drift-reduction ($\omega \ll \omega_{ci}$), and is due to the lower pressure gradient in the L-mode. In the H-mode, the drift-wave frequency reaches about half the cyclotron frequency at peak error, providing a stabilising effect to the drift-waves that does not exist in the full-velocity system. Note that the peak disagreement occurs at the same radial location as the peak growth rate.

The error in the frequencies exhibits the same qualitative behaviour as the growth rate error, but with nearly five times the accuracy. The maximum error in the frequency is 7% in H-mode compared to 34% for the growth rate. This good agreement is due to the relatively small value of the additional real terms in the full dispersion relation, equation 6.

4.2 Mode number dependence

Even when constrained by the JET data, the accuracy of the drift-reduction is quite sensitive to the ratio of the parallel and perpendicular wavelengths of the drift-instability. It is useful to normalise the parallel wave number to the major radius, $R_0 = 296$cm on JET. Holding the perpendicular wave number constant at $k_z = 10\text{cm}^{-1}$ [16] and scanning the parallel wave number across a reasonable domain $k_y R_0 \in [2.5, 25]$, the trend for the error is seen to

be roughly exponentially decreasing with $k_y R_0$, as shown in figure 7. Even within this somewhat small range of parallel wave numbers, there is a significant difference in maximum growth rates ranging from an unacceptable 800% error down to 3% error for the H-mode plasma. The normalised parallel wavenumber can be defined as $k_y R_0 = n/2\pi q$ where n is the toroidal mode number of the drift-wave and q is the safety factor. This indicates that high mode number drift-waves (and perhaps other instabilities as well) are represented accurately by the drift-reduced model while those with low mode number must be described with a full-velocity model.

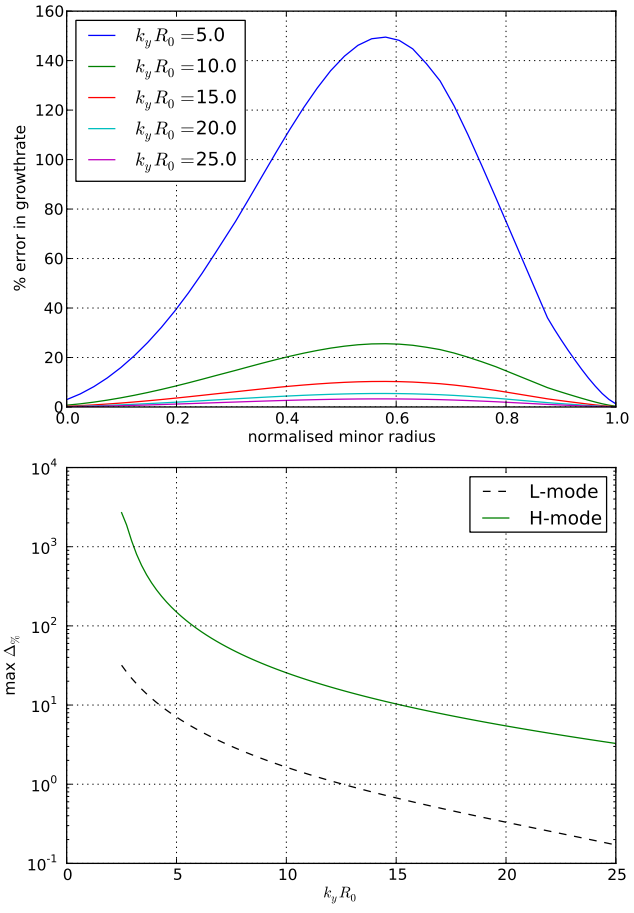

Figure 7: The percentage error in growth rate between the drift-reduced and full-velocity models is shown (for H-mode) to depend on the ratio of parallel to perpendicular wave number for the drift-wave instability (top). The maximum percentage error in growth rate is nearly an exponentially decaying function of parallel wave number (bottom).

5 Conclusion

Drift-reduced models provide simplified dispersion relations for more succinct analytics, and the exclusion of fast waves allows for larger time steps leading to faster simulations, so these models are an important subset of the full fluid description. The validity of these models has been tested for a quasi-3D slab resulting in drift-wave linear growth rates and frequencies that only agree with the full-velocity fluid description in specific regions of parameter space. Though the worst agreement lies outside of the operational regime of tokamaks, there is still poor agreement for the core plasma for reasonable plasma parameters, so one must use a full-velocity fluid model or instead employ a gyrokinetic description when modelling the core. Interestingly, the derivation of gyrokinetic models often employs the drift-reduction as well [17], so they may potentially suffer from similar inaccuracies at mid radius - a subject that requires further study.

When discussing the validity of drift-reduced models, it is necessary to consider the non-linear behaviour in addition to the linear. In the basic slab geometry investigated here, the linear differences are directly related to the non-linear saturated turbulent transport and growth times. The relationship between drift-wave linear and non-linear behaviour and mode structure in a more realistic *sheared* magnetic field is discussed in detail by Scott [18], and it is concluded that the non-adiabaticity of drift-waves is affected and usually enhanced by the non-linearities. This then drives the drift-waves further unstable, at which point the growth from the linear phase is irrelevant to the behaviour of the turbulence. That is not to say the linear growth rates do not play a role in the initial development of the turbulence. For example, in figure 3 it is apparent that for conductivities greater than $10^{13}\mathrm{s}^{-1}$ and electron beta, $\beta_e \in [10^{-2}, 10^{-3}]$, which corresponds to the tokamak core operating regime, the full velocity fluid model dictates that drift-waves are highly unstable whereas the drift-reduced model places them near marginal stability. It is not unreasonable, then, to assert that the linear and non-linear behaviours are highly correlated in this region, since in the drift-reduced case the modes can be easily stabilised preventing altogether the development of turbulence. Once the non-linear turbulence is established, however, it can be self-sustaining even if the linear modes are then sta-

bilised [18, 19].

It is important when choosing a fluid model to use for tokamak plasma simulations to identify the parameter space in which the simulation will be operating as to identify whether a drift-reduced model is appropriate or if a more accurate, full-velocity model should be used instead.

Acknowledgements

This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014-2018 under grant agreement No 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission. We thank Sarah Newton, Anantanarayanan Thyagaraja, and Arkaprava Bokshi for useful discussions guiding the evolution of this work and Alasdair Wynn for supplying the JET data. This work has received funding from the RCUK Energy Programme [grant number EP/I501045].

References

[1] S. Braginskii. Transport Processes in a Plasma. *Reviews of Plasma Physics*, 1:205–311, 1965.

[2] Bruce D Scott. Computation of turbulence in magnetically confined plasmas. *Plasma Physics and Controlled Fusion*, 48:B277–B293, 2006.

[3] E Frieman and L Chen. Nonlinear gyrokinetic equations for low-frequency electromagnetic waves in general plasma equilibria. *Physics of Fluids*, 25(1982):502–508, 1982.

[4] Gregory W. Hammett and Francis W. Perkins. Fluid moment models for Landau damping with application to the ion-temperature-gradient instability. *Physical Review Letters*, 64(25):3019–3022, 1990.

[5] Bruce D. Scott. Drift wave versus interchange turbulence in tokamak geometry: Linear versus nonlinear mode structure. *Physics of Plasmas*, 12(6):062314, 2005.

[6] V P Pastukhov, N V Chudin, and D V Smirnov. Effective fluid model of turbulent dynamics and transport in tokamak core plasmas. *Plasma Physics and Controlled Fusion*, 53:054015, 2011.

[7] A. B. Mikhailovskii and V. S. Tsypin. Transport equations and gradient instabilities in a high pressure collisional plasma. *Plasma Physics*, 13(9), 1971.

[8] A. N. Simakov and P. J. Catto. Drift-ordered fluid equations for modelling collisional edge plasma. *Contributions to Plasma Physics*, 44(1):83–94, 2004.

[9] P. H. Rutherford. Drift Instabilities in General Magnetic Field Configurations. *Physics of Fluids*, 11(1968):569, 1968.

[10] V. Naulin and K. H. Spatschek. Nonlinear drift-wave structures and their influence on particle transport. *Physical Review E*, 55(5):5883–5893, 1997.

[11] P.J. Knight, A. Thyagaraja, T.D. Edwards, J. Hein, M. Romanelli, and K.G. McClements. CENTORI: A global toroidal electromagnetic two-fluid plasma turbulence code. *Computer Physics Communications*, 183(11):2346–2363, November 2012.

[12] R. D. Hazeltine, M. Kotschenreuther, and P. J. Morrison. A four-field model for tokamak plasma dynamics. *Physics of Fluids*, 28(8):2466, 1985.

[13] E M Lifshitz and L P Pitaevskii. *Course of Theoretical Physics: Physical Kinetics*, volume 10. 1981.

[14] J. Wesson and D.J. Campbell. *Tokamaks*. International Series of Monographs on Physics. OUP Oxford, 2011.

[15] L.L. Lao, H. John, R.D. Stambaugh, a.G. Kellman, and W. Pfeiffer. Reconstruction of current profile parameters and plasma shapes in tokamaks. *Nuclear Fusion*, 25(11):1611–1622, 1985.

[16] D. L. Brower, W. A. Peebles, N. C. Luhmann, and R. L. Savage. Multichannel Scattering Studies of the Spectra and Spatial Distribution of Tokamak Microturbulence. *Physical Review Letters*, 54(7):689–692, 1985.

[17] A. J. Brizard and T. S. Hahm. Foundations of nonlinear gyrokinetic theory. *Reviews of Modern Physics*, 79(June):421–468, 2007.

[18] B. D. Scott. Self-Sustained Collisional Drift-Wave Turbulence in a Sheared Magnetic Field. *Physical Review Letters*, 65(26), 1990.

[19] Hans Nordman, Vladimir P. Pavlenko, and Jan Weiland. Subcritical reactive drift wave turbulence. *Physics of Fluids B: Plasma Physics*, 5:402, 1993.