Asymptotic and non-asymptotic model reduction for kinetic descriptions of plasma

John B. Coughlin

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2024

Reading committee:

Uri Shumlak, Chair

Jingwei Hu, Chair

Eric Meier

Program Authorized to Offer Degree:
Applied Mathematics

University of Washington

Abstract

Asymptotic and non-asymptotic model reduction for kinetic descriptions of plasma

John B. Coughlin

Chairs of the Supervisory Committee:

Uri Shumlak
Department of Aeronautics & Astronautics

Jingwei Hu
Department of Applied Mathematics

Plasma dynamics are coupled across microscopic and macroscopic scales by a variety of nonlinear mechanisms. These include repartition of energy due to kinetic microinstabilities, suppression of fluid instabilities by kinetic stabilization effects, and other mechanisms. At the macroscopic scale plasmas are well-described by fluid equations, whose formal validity depends on the long-time regularization of the phase space distribution function by collisions and magnetic gyrotropization. However, accurately capturing multiscale coupling requires multiscale reduced models which are both efficient and accurate in transition regimes. These regimes, where either collisional or magnetic gyrotropic regularization are marginal, are characterized by the ratio of the (collisional or magnetic) mean free path to a characteristic gradient scale length. This work studies two families of reduced plasma models for transition regimes in depth. The first is an asymptotic expansion for Braginskii-type transport coefficients in the so-called drift ordering for low-beta plasmas. The expansion captures leading-order finite Larmor radius effects for arbitrary collisionality. We present a new derivation of this expansion, evaluate its performance numerically, and provide a numerically feasible approximation. The second family of methods is dynamical low-rank (DLR) methods, which are not based on an asymptotic expansion and have the potential to overcome the curse of dimensionality for kinetic equations. We present two novel DLR schemes for plasma kinetic equations with a focus on fluid-kinetic coupling. One is a DLR method that retains low rank in the highly collisional asymptotic limit. The other is a fully locally conservative DLR method for the Vlasov-Dougherty-Fokker-Planck equation which achieves second-order accuracy in time. All discretizations are described in detail and accompanied by numerical results demonstrating the merit of the proposed approach.

Dedication

To my wife Lizeth

Acknowledgements

Thank you to my indefatigable, ever-patient and ever-curious co-advisor Uri Shumlak. From our very first meeting you have always projected the same attitude about this ridiculous challenge we set ourselves to develop nuclear fusion energy: *if you've got any ideas I'd love to hear them.* I can't overstate how beneficial this steadfast open-mindedness is to the sort of young researcher I was five years ago. Thank you for your sense of humor and most of all for the incredible level of support you worked to provide for me in the course of my long doctoral studies.

Thank you to my energetic, many-talented co-advisor Jingwei Hu. Your willingness to jump immediately into working together on plasma problems helped me immensely during a time when I was struggling to make concrete progress. You have a wealth of knowledge about deep technical topics, but you somehow always made me feel treated as a peer, even when asking the most basic questions. Thank you for the mathemetical community you have introduced me to, for the sophistication that you insisted on in our papers, and for all your support.

Thank you also to the UW Computational Plasma Dynamics group—to Eric Meier, Whitney Thomas, Iman Datta, Daniel Crews, Adrian Carillo, Yu Takagaki, Evan Bluhm, Howard Cheng, Noah Reddell, Andrew Ho, and Gabriel Rodriguez. The thoughtful discussions we had have been enormously helpful during the development of the ideas in this thesis. And at the risk of repeating myself, thank you to the Theory and Modeling team at Zap Energy, both for a challenging internship and for the opportunity to return for more.

I would not have been able to write any of this thesis if it were not for the excellent training I received from the Applied Mathematics department. Thank you to all of the professors whose classes I took, especially Professor Emeritus Randall Leveque, whose red book is a permanent fixture on my cluttered desk. Thank you also to Professor Justin Little of the Aerospace and Astronautics department for his excellent class on plasma kinetic theory.

Thank you to all of the brilliant researchers who gave advice, asked questions, offered opportunities to present, and were good friends at conferences—thank you to Genia Vogman, Lukas Einkemmer, Jingmei Qiu, Andrew Christlieb, Luis Chacon, Jimmy Juno, Joseph Nakao, and others.

Finally, a million thanks to my family, to Scott, Lisa, Jamie, for your support, friendship, and curiosity. Thank you for wanting to come to all of my talks. Thank you dad for all of the articles. Most of all, thank you to my wife Lizeth, without whom this thesis would never have been possible. I can't imagine a more supportive life partner. Thank you for grounding me and for always being there for me. I can't wait to return that support to you during our next chapter, whatever it is. And thank you to our dog Musa for showing us your blanket every morning, we couldn't have done it without you.

Funding Acknowledgement

The information, data, or work presented herein is based upon work supported by the National Science Foundation under Grant No. PHY-2108419.

Contents

Chapter 1

Introduction

The overwhelming majority of the baryonic matter in the universe, including every main sequence star and most of the stuff in between them, is in the plasma state. On Earth, plasmas can be found in lightning bolts, extreme ultraviolet lithography machines, and nuclear fusion reactors. Across all of these arenas, the same equations describe physical phenomena that span 20 orders of magnitude in time and length scales. The 10 ns-long National Ignition Facility shot N210808 [1] is described by the same fundamental equations as the million-year evolution of an astrophysical plasma jet.

Plasmas consist of charged particles—typically ions and electrons—which at different scales appear to behave variously like a single fluid, or like multiple separate interacting fluids, or, at the smallest scales, unlike anything else in nature. Zooming out quite far, we find that the single fluid *magnetohydrodynamic* model is accurate and efficient. At smaller scales, ions have enough inertia that they flow ahead of electrons in the directions perpendicular to the magnetic field and lag behind them in the parallel direction, introducing charge separation physics that requires a *multifluid* model to capture. And at the smallest scales, from the perspective of the electrons, the lumbering ions appear stationary, so that the electrons get up to esoteric *kinetic* physics right under their nose.

This picture of different models, each behaving nicely at their respective scale, is unfortunately far too neat, for the reason that interesting physics cascades from large scales down to small, and likewise filters back up from small scales to large. For the purposes of introduction we mention two examples. Both have to do with the concept of instability, which is a primary mechanism for small-scale physics to magnify into large-scale effects. The first is a kinetic phenomenon: the Dory-Guest-Harris instability [29], which can arise from unstable loss-cone velocity distributions in magnetic mirror confinement, has the effect of turning free thermal energy into electrostatic wave energy. Particles destabilized in this way can re-enter the loss cone and exit confinement. The second example is also a kinetic effect, but one which manifests as a kind of fluid viscosity. In magnetized plasmas, communication of momentum across the radius of a particle's cyclotron orbit results in a transport phenomenon called gyroviscosity, which sets a lower limit on the gradient scale lengths that the plasma can exhibit. This in turn affects the growth rate of fluid instabilities such as the Rayleigh-Taylor instability [63, 65, 66]. Thus, the details of a kinetic particle distribution function can have a large impact on observable physics at the fluid scale.

From a mathematical perspective, the analysis of plasmas and their governing equations begins with the Vlasov-Boltzmann equation,

$$\partial_t f_a + \boldsymbol{v} \cdot \nabla f_a + \frac{q_a}{m_a}(\boldsymbol{E} + \boldsymbol{v} \times \boldsymbol{B}) \cdot \nabla_{\boldsymbol{v}} f_a = C(f). \tag{1.0.1}$$

The scalar function $f(\boldsymbol{x}, \boldsymbol{v}, t)$ depends on six phase-space coordinates and time, so the direct solution of (1.0.1) is numerically very challenging. There are two principal effects which make solutions of (1.0.1) well-behaved, in the sense that their behavior can be described by lower-dimensional fluid models. These are:

- **Magnetization**: If the magnetic field $\boldsymbol{B}$ is very strong, the plasma behavior in the directions perpendicular to $\boldsymbol{B}$ can be described by fluid-like equations.
- **Collisionality**: If the effect of collisions $C(f)$ is very strong, the behavior of each plasma species will be fluid-like.

We stress that whether the magnetization or collisionality is "strong" is as much a property of the chosen normalization (the "zoom level") as it is of the physical situation. The normalization is prescribed in terms of the fundamental lengths and frequencies of the plasma: the Debye length, cyclotron frequency, plasma frequency, etc. For example, the two-fluid instability takes place on the order of a few dozen Langmuir oscillation periods, and so the corresponding normalized Vlasov equation can neglect the much slower collision frequency completely. On the other hand, if the time period of interest spans an entire five-second tokamak shot, collisions will be much "stronger" in the normalized equation, simply because they have more time to act.

If a plasma is highly magnetized or highly collisional, the benefits from using a fluid model instead of a kinetic model can be immense. Fluid models evolve unknowns that depend on $\boldsymbol{x}$ and t, but not on $\boldsymbol{v}$: the velocity dependence is integrated away. For three-dimensional problems, a fluid model can easily have a factor of 10^6 fewer degrees of freedom than a kinetic model for the same problem. While the development of optimal numerical methods for plasma fluid models is still an ongoing area of research, it is fair to say that the underlying equations describing highly magnetized and highly collisional plasmas are very well understood.

In this thesis we describe several contributions to the theory and numerical simulation of plasmas which are *marginally* magnetized or collisional. By marginally magnetized or marginally collisional, we mean plasmas which exhibit some non-ideal or kinetic effect which cannot be neglected. These effects can include collision-induced viscosity and resistance in multispecies plasmas, finite Larmor radius (FLR) effects such as gyroviscosity in magnetized plasmas, and free-streaming physics such as heat diffusion.

A core theme of this work is the dichotomy between the asymptotic strategy for model reduction and non-asymptotic strategies. By asymptotic model reduction we refer to all approaches based on arguing that a given term in (1.0.1) is very large or very small. The asymptotic methodology introduces a small parameter ϵ and performs series expansions of the kinetic equation, thereby deriving zeroth-order solutions plus first-order and higher corrections. The resulting theories can be very powerful and predictive, but by design they have limited regimes of applicability: if ϵ is no longer small enough that it or its square can be ignored, the asymptotic theory no longer applies. A benefit of asymptotic theories is that they typically result in reduced models that preserve important structure from the parent kinetic model. For example, the Braginskii transport theory [11] is naturally conservative, with diffusive terms that each correspond to recognizable physics: resistivity, heat conduction, and viscosity. However, the multiscale nature of plasmas, outlined above, means that non-asymptotic model reduction methods which can apply across a range of scales have the potential to greatly improve our simulation capabilities.

In line with this framing, we introduce our three major contributions in this thesis

- A new derivation and numerical evaluation of the so-called “drift ordering” asymptotic expansion for collisionless, magnetized plasmas.
- A pair of structure-preserving methods based on the dynamical low-rank (DLR) methodology [71], including
  - A DLR algorithm for the Vlasov-Ampère-Fokker-Planck equation with a linearized collision operator which retains the efficiency improvements of solving a fluid equation in the highly collisional regime [19].
  - A DLR algorithm for the Vlasov-Dougherty-Fokker-Planck equation that conserves charge, momentum and energy exactly [20].

We introduce these in somewhat more detail in the following sections.

Asymptotic transport theory of low-beta collisionless plasmas

The famous Braginskii [11] transport theory is a highly successful classical asymptotic expansion. It describes certain collisional and FLR effects in highly collisional, magnetized plasmas. It is based on an expansion in the Knudsen number, which is defined as the ratio of the collisional mean free path to the characteristic scale length. Unfortunately, the assumption of strong collisionality is inconvenient for many plasma regimes. Based on this, Z-Pinch [81, 98] and tokamak [42] modeling efforts have employed correction factors to the Braginskii coefficients. These correction factors attempt to account for the fact that collisional energy equipartition between parallel and perpendicular velocities may be quite slow in reality. Additionally, [81] presented correction factors to partially account for FLR effects in Z pinches.

By assuming strong collisions, the Braginskii derivation is able to proceed without making any assumption on the magnitude of the fluid flow velocity. On the other hand, relaxing the collisional assumption still allows for an asymptotic expansion around a gyrotropic, and thus necessarily stationary, Maxwellian distribution. The leading-order perpendicular velocity is then given by the $E \times B$ and diamagnetic drifts. Such an ordering has been used before to derive reduced models for magnetized plasmas, including reduced magnetohydrodynamics [57, 61] and multi-fluid models [86, 87].

Chapter 3 presents a new derivation of the drift ordering closure for the Vlasov equation, based on less restrictive assumptions than [11]. In line with prior work on asymptotics of the collisionless Vlasov equation, we make no assumption of strong collisionality. We assume strong magnetization (or low beta), and that the marginal distribution function f in the perpendicular direction is nearly Maxwellian. These assumptions let us derive asymptotic closures for a set of “semi-fluid” equations, which are kinetic in the direction parallel to the magnetic field, and fluid perpendicular to it. The resulting transport relations constitute corrections to the Braginskii transport relations for this regime. In particular, we find that if all fluid quantities are expanded consistently to second-order in the expansion parameter, the gyroviscous stress contains non-negligible terms related to the heat flux. The modified gyroviscosity closure is directly applicable to kinetic-fluid multiscale coupling such as the Rayleigh-Taylor instability example mentioned above.

Chapters 4 and 5 contain a numerical validation of the claims made in Chapter 3. Chapter 4 describes a Fourier-Hermite spectral Vlasov solver specifically designed for solving low-beta plasma flow problems to high accuracy. The code makes use of a pseudospectral Fourier expansion in periodic dimensions in $\boldsymbol{x}$, and a global Hermite spectral discretization of velocity space. The Hermite discretization gives excellent resolution of distribution functions in the weakly compressible flow regime. The code is also equipped with a high-order Shu-Osher form finite difference discretization in one spatial dimension, enabling us to solve non-periodic problems involving flow shear induced by nonzero charge density.

In Chapter 5, we apply the spectral Vlasov solver to a family of initial conditions exhibiting density and temperature gradients as well as sheared flow. We employ a simple statistical methodology to analyze the validity of drift ordering transport closure compared to the classical Braginskii closure. Numerical results indicate that the drift ordering closure more accurately predicts the magnitude of the gyroviscous stress for all configurations tested, and that the modification compared to the Braginskii closure is significant at leading order in ϵ. Analysis of residuals shows that second-order and higher effects play a significant role in determining the spatial structure of heat flux and gyroviscous stress. To facilitate the numerical implementation of the drift ordering gyroviscous stress closure, we propose a numerically affordable approximation. The structure of the approximate form allows us to predict that the importance of the heat flux correction is connected to the ratio of diamagnetic to $E \times B$ drift velocity: larger $E \times B$ drifts correspond to reduced significance of the diamagnetic drift and heat flux.

Dynamical low-rank methods for weakly collisional plasmas

In addition to asymptotic model reduction techniques, the multiscale challenges of plasma dynamics motivate the development of non-asymptotic techniques, which rely on some other assumed property of the kinetic equation to solve it in an approximate but hopefully much more efficient manner. Perhaps controversially, we mention Particle-in-Cell (PIC) methods [9] as an example. By approximating f_s as a superposition of δ functions which act as macro-particles, PIC methods trade the high dimensionality of phase space for statistical noise. PIC methods have seen enormous success in many areas of computational plasma physics.

Chapters 6-8 describe our work on another family of non-asymptotic model reduction techniques for the Vlasov equation. These are dynamical low-rank (DLR) methods [72]. DLR methods are based on the assumption that the data contained in the phase space distribution is highly compressible in the sense of having a low numerical rank when viewed as a matrix or tensor. This low-rank *ansatz* is of dubious physical provenance, but has been found to be quite successful in practice [36, 32] when applied to kinetic plasma problems. However, preserving important structure such as asymptotic limits or conservation laws is a challenge when using DLR methods.

Therefore, Chapter 7 describes an efficient DLR algorithm for a simple collisonal model equation which preserves the Maxwellian asymptotic limit with a low rank requirement. The algorithm is based on applying a low-rank *ansatz* to the quotient $g = f/M$ where M is the local Maxwellian. Chapter 8 presents a locally conservative DLR scheme for the Vlasov-Fokker-Planck equation based on a novel macro-micro decomposition of f. Whereas previous conservative DLR schemes required a modification to the DLR equations of motion [34, 38], our macro-micro decomposition naturally separates conserved from non-conserved quantities at the equation level. The result is a continuous system of equations which can be flexibly integrated by any DLR integrator, and which can be discretized using a variety of schemes in velocity space. This flexibility is important for developing future tensor-based schemes for higher-dimensional plasma problems and for resolving filamentation in phase space.

Organization of the thesis

This work is organized as follows. Chapter 2 contains precise statements of the plasma models we consider here, including the Vlasov-Boltzmann equation, five-moment fluid equations, and Maxwell's equations of electromagnetism. It also contains a description of the normalization scheme used in the rest of the work. Chapter 3 derives the leading-order transport theory for the collisionless Vlasov equation in the strongly magnetized drift ordering. Chapter 4 describes a Fourier-Hermite spectral discretization of the Vlasov equation, which is used in Chapter 5 to numerically compare the performance of the asymptotic moment closures of Braginskii and Chapter 3.

Chapters 6-8 are devoted to the dynamical low-rank (DLR) family of methods. Chapter 6 provides background and theoretical grounding for the DLR method in the abstract setting, as well as an illustration of how it may be applied to the Vlasov equation. Chapter 7 is drawn from our paper [19], and contains a description of an efficient DLR algorithm for the collisional Vlasov-Ampère-Fokker-Planck equation. Chapter 8 reproduces work published in [20], and describes a novel macro-micro decomposition for the Vlasov equation which leads to a DLR scheme satisfying exact discrete local conservation laws. Chapter 9 concludes by summarizing the research presented, discussing future directions, and connecting to broader themes in computational plasma physics.

Chapter 2

Mathematical models of plasma

This chapter describes several mathematical models of plasma at different levels of fidelity. When deriving a lower-fidelity model from a higher-fidelity one, we typically seek to trade the ability to capture certain physical effects for greater analytic and computational tractability. It is hoped that the low-fidelity model captures those physical effects which dominate in a given dynamical or equilibrium situation of interest.

This chapter begins with a presentation of the Vlasov equation, including its collisional variants, the Vlasov-Boltzmann, Vlasov-Landau, and Vlasov-Fokker-Planck equations. We then derive the exact five-moment fluid equations and describe the moment closure problem, with the ideal five-moment fluid equations as a specific choice of closure. Section 2.3 presents Maxwell's equations, which are of key importance in plasma physics. Finally, Section 2.4 discusses the flexible plasma normalization used in this thesis.

2.1 Kinetic models: the Vlasov equation and its siblings

The core object of interest in plasma kinetic theory is invariably denoted f_a. It is the *particle distribution function* for particles of species a:

$$f_a(\boldsymbol{x}, \boldsymbol{v}, t) : \Omega_{\boldsymbol{x}} \times \mathbb{R}^d \times \mathbb{R} \mapsto \mathbb{R}.$$

Here, $\Omega_{\boldsymbol{x}}$ is a domain in physical space and d is the number of independent velocity dimensions considered. The value of f_a represents the density of particles at position $\boldsymbol{x}$ and time t with velocity $\boldsymbol{v}$. That is, we always use the normalization

$$n_a(\boldsymbol{x}, t) = \int_{\mathbb{R}^d} f_a \,\mathrm{d}\boldsymbol{v} \tag{2.1.1}$$

where n_a is the number density of species a, having dimensions of particles per volume.

The particle distribution f_a evolves in time according to the Vlasov-Boltzmann equation,

$$\partial_t f_a + \boldsymbol{v} \cdot \nabla_{\boldsymbol{x}} f_a + \frac{q_a}{m_a}(\boldsymbol{E} + \boldsymbol{v} \times \boldsymbol{B}) \cdot \nabla_{\boldsymbol{v}} f_a = \left(\frac{\partial f_a}{\partial t}\right)_c, \tag{2.1.2}$$

where $\left(\frac{\partial f_a}{\partial t}\right)_c$ represents the effect of collisions. Equation (2.1.2) must be coupled to Maxwell's equations (see Section 2.3) for the electric field $\boldsymbol{E}$ and magnetic field $\boldsymbol{B}$.

The Vlasov equation is the kinetic equation which takes the right-hand side to be zero:

$$\partial_t f_a + \boldsymbol{v} \cdot \nabla_{\boldsymbol{x}} f_a + \frac{q_a}{m_a}(\boldsymbol{E} + \boldsymbol{v} \times \boldsymbol{B}) \cdot \nabla_{\boldsymbol{v}} f_a = 0. \tag{2.1.3}$$

This equation is widely used in plasma physics due to the extremely low frequency of collisions in many plasmas of interest. It is worth remarking briefly on the nature of the Vlasov equation, which is a time-dependent hyperbolic advection equation in six-dimensional phase space. It contains two advection terms:

- the "free-streaming" term $\boldsymbol{v} \cdot \nabla_{\boldsymbol{x}} f_a$, which says that a chunk of particle density moves through physical space with its own velocity $\boldsymbol{v}$, and
- the Lorentz force term, which says that a chunk of particle density is accelerated by a force equal to the Lorentz force $q_a(\boldsymbol{E} + \boldsymbol{v} \times \boldsymbol{B})$.

In this thesis we will use ∇ in place of $\nabla_{\boldsymbol{x}}$ whenever it is unambiguous.

2.1.1 Binary collisions

The purpose of the collision term, which we have denoted $\left(\frac{\partial f_a}{\partial t}\right)_c$, is to capture the effect on the distribution function f_a of collisions between particles of species a and particles of all other species (including a). A key assumption of collisional kinetic modeling is that collisions are *binary* [14], and that we may write

$$\left(\frac{\partial f_a}{\partial t}\right)_c = \sum_b C_{ab}(f_a, f_b), \tag{2.1.4}$$

for some binary collision operator C_{ab} representing the action of elastic collisions between particle species a and b on the distribution function f_a. Of course, there must be a corresponding operator C_{ba} which acts on f_b and averages over the same collision events.

Collisional processes in plasmas are quite distinct from those in neutral gases for the following reason: the Coulomb interaction strength falls off as $1/r^2$ where r is the distance between interacting particles. For this reason, the vast majority of collisions in plasmas are "grazing", i.e. small-angle collisions, in which two particles scatter off each other at a great distance, changing their trajectories only a small amount. This has implications for the treatment of collisions in the continuum kinetic picture. The mathematical model which is widely accepted as physically accurate for two-particle Coulomb collisions is the Landau collision operator [75]:

$$C_{ab}(f_a, f_b)(\boldsymbol{v}) = \frac{\gamma_{ab}}{m_a} \nabla_{\boldsymbol{v}} \cdot \int_{\mathbb{R}^3} \mathbf{A}(\boldsymbol{v} - \boldsymbol{v}_*) \left[\nabla_{\boldsymbol{v}} f_a \frac{f_b}{m_a} - \frac{f_a}{m_b} \nabla_{\boldsymbol{v}_*} f_b \right] \mathrm{d}\boldsymbol{v}_*, \tag{2.1.5}$$

where the kernel $\mathbf{A}$ is a 3×3 tensor

$$\mathbf{A}(\boldsymbol{z}) = \frac{|\boldsymbol{z}|^2 \mathbb{I} - \boldsymbol{z} \otimes \boldsymbol{z}}{|\boldsymbol{z}|^3}$$

and $\gamma_{ab} = 2\pi \ln \Lambda q_a^2 q_b^2 / 8\pi\epsilon_0^2$, with $\ln \Lambda \approx 10$ the Coulomb logarithm [40]. The operator (2.1.5) is derived from the more general Boltzmann operator by retaining terms up to second order in the difference between pre- and post-collision velocities.

The Landau collision operator is an integro-differential operator, whose evaluation requires accounting for contributions that are non-local in velocity space. Far more efficient numerical methods can be obtained if one makes further simplifications to obtain an operator that is local in velocity space. However, all useful model collision operators must respect certain key properties of the Landau (and Boltzmann) operators, in particular conservation and the so-called H theorem, which we describe below.

Regarding conservation, recall that the collision operator C_{ab} models elastic collisions, meaning collisions that preserve the total momentum and energy of the two participant particles. It is possible to consider additional collision terms that introduce reactions, including particle fusion and creation, but we neglect these effects here. Mathematically, the requirement that binary collisions preserve particle mass, momentum and energy can be expressed in terms of velocity moments of the collision operator:

$$\int_{\mathbb{R}^d} C_{ab}(f_a, f_b) \, \mathrm{d}\boldsymbol{v} = 0 \tag{2.1.6}$$

$$\int_{\mathbb{R}^d} m_a \boldsymbol{v} C_{ab}(f_a, f_b) \, \mathrm{d}\boldsymbol{v} + \int_{\mathbb{R}^d} m_b \boldsymbol{v} C_{ba}(f_b, f_a) \, \mathrm{d}\boldsymbol{v} = 0 \tag{2.1.7}$$

$$\int_{\mathbb{R}^d} \frac{m_a |\boldsymbol{v}|^2}{2} C_{ab}(f_a, f_b) \, \mathrm{d}\boldsymbol{v} + \int_{\mathbb{R}^d} \frac{m_b |\boldsymbol{v}|^2}{2} C_{ba}(f_b, f_a) \, \mathrm{d}\boldsymbol{v} = 0. \tag{2.1.8}$$

In addition to conservation, another key property of inter-particle collisions is that they increase entropy. Consider the spatially homogeneous Boltzmann equation for a neutral species a,

$$\partial_t f_a = \sum_b C_{ab}(f_a, f_b). \tag{2.1.9}$$

It can be shown [106] that if each distribution function f_a satisfies (2.1.9), then the following inequality holds:

$$\frac{\mathrm{d}}{\mathrm{d}t} \sum_a H(f_a) \triangleq \frac{\mathrm{d}}{\mathrm{d}t} \sum_a \int_{\mathbb{R}^d} f_a \log f_a \, \mathrm{d}\boldsymbol{v} \leq 0.$$

This is the celebrated H-theorem of Boltzmann, and it provides a useful and physical constraint on simplified models for collision operators. In particular, the conservation properties (2.1.6) and the H-theorem imply [14, 106] that the kernel of C_{ab} is the Maxwellian, which in d velocity space dimensions has the form

$$\mathcal{M}_a(\boldsymbol{x},\boldsymbol{v},t) = n_a(\boldsymbol{x})\left(\frac{m_a}{2\pi T}\right)^{d/2} e^{-\frac{m_a|\boldsymbol{v}-\boldsymbol{u}(\boldsymbol{x})|^2}{2T(\boldsymbol{x})}}, \tag{2.1.10}$$

where the moments $\boldsymbol{u}$ and T can be determined from momentum and energy conservation of the multispecies particle system.

The existence of an equilibrium solution to (2.1.9) immediately suggests a simple model collision operator, where the distribution function is taken to relax uniformly and linearly to its equilibrium. The Bhatnagar-Gross-Krook (BGK) operator [8, 49] between species a and b is

$$C_{ab}^{BGK}(f_a, f_b) = \nu_{ab}(\mathcal{M}_{ab} - f_a), \tag{2.1.11}$$

where ν_{ab} is a collision frequency and $\mathcal{M}_{ab}$ is the Maxwellian with average velocity $\boldsymbol{u}_{ab}$ and temperature T_{ab}. Various models for the collision moments $\boldsymbol{u}_{ab}$ and T_{ab} have been proposed. Pairwise conservation constrains but does not determine the collision moments, and the collision frequency ν_{ab} arises as a free parameter which can be chosen so that the BGK operator mimics certain relaxation rates of the full Landau operator. We refer to [49] for an overview of the work in this area.

The BGK model is extremely useful and simple to implement. However, in the plasma context, it neglects the essentially diffusive character of the Landau collision operator. Another widely used model collision operator is the Lenard-Bernstein (LBO) collision operator [76], which was also studied by Dougherty [30]. This has the form

$$C_{ab}^{LBO}(f_a, f_b) = \nu_{ab}\nabla_{\boldsymbol{v}} \cdot \left(\frac{T_{ab}}{m_a}\nabla_{\boldsymbol{v}} f_a + (\boldsymbol{v} - \boldsymbol{u}_{ab}) f_a\right). \tag{2.1.12}$$

Evidently this is constructed as a Fokker-Planck type operator, with diffusion coefficient T_{ab}/m_a and drift velocity $\boldsymbol{u}_{ab}$. The quantities T_{ab} and $\boldsymbol{u}_{ab}$ can be thought of as the temperature and average velocity of a "target" Maxwellian towards which C_{ab}^{LBO} is driving f_a. Indeed, it is easy to show that the Maxwellian with those parameters is in the kernel of C_{ab}^{LBO}. These "collisional moments" are constrained by the requirement of conservation of momentum and energy in binary collisions between species a and b. Substituting (2.1.12) into (2.1.6) and integrating by parts, we obtain the following pair of expressions:

$$\boldsymbol{u}_{ab} = \frac{\nu_{ab} m_a n_a \boldsymbol{u}_a + \nu_{ba} m_b n_b \boldsymbol{u}_b}{\nu_{ab} m_a n_a + \nu_{ba} m_b n_b} \tag{2.1.13}$$

$$T_{ab} = \frac{T_a \nu_{ab} n_a + T_b \nu_{ba} n_b}{\nu_{ab} n_a + \nu_{ba} n_b} - \frac{\nu_{ab} n_a m_a (\boldsymbol{u}_{ab}^2 - \boldsymbol{u}_a^2) + \nu_{ba} n_b m_b (\boldsymbol{u}_{ba}^2 - \boldsymbol{u}_b^2)}{3(\nu_{ab} n_a + \nu_{ba} n_b)}, \tag{2.1.14}$$

where the moments $n_a, \boldsymbol{u}_a$ and T_a are defined by (2.1.1), (2.2.2), and (2.2.4). In the derivation we have made the natural pair of assumptions $\boldsymbol{u}_{ab} = \boldsymbol{u}_{ba}$ and $T_{ab} = T_{ba}$, which, in physical terms, mean that particles of different species are driven towards the *same* Maxwellians, differing only in the velocity scale attributable to the mass ratio m_a/m_b.

It remains only to determine the collision frequency ν_{ab}. This can be done [49, 100] by requiring that the rate of relaxation of momentum due to collisions match the rate obtained from a more physical collision operator, in this case the Landau collision operator. The choice made here corresponds to the $\nu_{ab}^{EM,LBO}$ collision frequency as used in [100]. Converting equation (13) in that reference from Gaussian to SI units,

$$\nu_{ab} = \nu_{ab}^{EM,LBO} = n_b \frac{\sqrt{m_a m_b}(m_a + m_b)(q_a q_b)^2 \ln \Lambda}{3\sqrt{2}\pi^{3/2}\epsilon_0^2 m_a (m_a T_b + m_b T_a)^{3/2}}. \tag{2.1.15}$$

Finally, we cite a recent development for the multispecies LBO collision moments. In [50] it is shown that by relaxing interspecies symmetry by allowing $\boldsymbol{u}_{ab} \neq \boldsymbol{u}_{ba}$ and $T_{ab} \neq T_{ba}$, both the momentum and temperature relaxation rates can be matched to the full Landau collision operator. This is a very promising result, in that preserving the physical momentum and temperature relaxation rates seems more important than the natural-seeming but restrictive symmetry of cross-species moments.

2.2 Five-moment fluid models

To derive simpler models of plasma than the Vlasov or Vlasov-Boltzmann equations, we integrate away the dependence on velocity to obtain (nonlinear) systems of equations posed over physical space alone. The five-moment model takes moments of f_a corresponding to the collision invariants $(1, \boldsymbol{v}, |\boldsymbol{v}|^2/2)^T$, to obtain a continuity equation, a momentum equation, and an energy equation.

Recalling the definition of n_a in (2.1.1), by taking the moment of (2.1.2) with respect to 1, we get

$$\partial_t n_a + \nabla \cdot (n_a \boldsymbol{u}_a) = 0, \tag{2.2.1}$$

where $\boldsymbol{u}_a$ is defined by

$$n_a \boldsymbol{u}_a = \int \boldsymbol{v} f_a \, \mathrm{d}\boldsymbol{v}. \tag{2.2.2}$$

Equation (2.2.1) is the continuity equation for species a.

We take the moment of (2.1.2) with respect to $m_a \boldsymbol{v}$ to obtain

$$m_a \partial_t (n_a \boldsymbol{u}_a) + \nabla \cdot (m_a n_a \boldsymbol{u}_a \otimes \boldsymbol{u}_a + p_a \mathbb{I} + \Pi_a) = n_a q_a (\boldsymbol{E} + \boldsymbol{u}_a \times \boldsymbol{B}) + \sum_b \boldsymbol{S}_{ab}. \tag{2.2.3}$$

Here we have introduced the scalar pressure

$$p_a = n_a T_a,$$

where T_a is defined by

$$T_a = \frac{m_a}{d n_a} \int_{\mathbb{R}^d} |\boldsymbol{v} - \boldsymbol{u}_a|^2 f_a \,\mathrm{d}\boldsymbol{v} \tag{2.2.4}$$

for a d-dimensional velocity space. The scalar pressure is proportional to the trace of the full pressure tensor $\mathbb{P}_a$, which is composed of the scalar pressure and a trace-free part Π_a:

$$\mathbb{P}_a = p_a \mathbb{I} + \Pi_a = m_a \int_{\mathbb{R}^d} (\boldsymbol{v} - \boldsymbol{u}_a) \otimes (\boldsymbol{v} - \boldsymbol{u}_a) f_a \,\mathrm{d}\boldsymbol{v}$$

where $\mathbb{I}$ is the $d \times d$ identity tensor. The tensor Π_a is called the stress tensor, and represents the non-isotropic, or "deviatoric", part of the full pressure tensor $\mathbb{P}_a$. We have also introduced the frictional forces $\boldsymbol{S}_{ab}$, which represent the transfer of momentum from species b to species a due to collisions, defined

$$\boldsymbol{S}_{ab} = m_a \int_{\mathbb{R}^d} \boldsymbol{v} C_{ab}(f_a, f_b) \,\mathrm{d}\boldsymbol{v}. \tag{2.2.5}$$

Finally, taking the moment of (2.1.2) with respect to $m_a \frac{|\boldsymbol{v}|^2}{2}$, we get the energy equation:

$$\partial_t e_a + \nabla \cdot (e_a \boldsymbol{u}_a + \boldsymbol{u}_a \cdot \mathbb{P}_a + \boldsymbol{q}_a) = n_a q_a \boldsymbol{u}_a \cdot \boldsymbol{E}_a + \sum_b Q_{ab}. \tag{2.2.6}$$

The quantity e_a is the total energy density of species a, which consists of kinetic energy and internal energy:

$$e_a = \frac{m_a n_a |\boldsymbol{u}_a|^2}{2} + \frac{p_a}{(\gamma - 1)}, \tag{2.2.7}$$

where for the particles in a plasma we use $\gamma = \frac{d+2}{d}$, indicating no internal degrees of freedom. In practical calculations, (2.2.7) is inverted to compute p_a. The vector $\boldsymbol{q}_a$ is the heat flux vector, defined as

$$\boldsymbol{q}_a = \frac{m_a}{2} \int_{\mathbb{R}^d} (\boldsymbol{v} - \boldsymbol{u}_a) |\boldsymbol{v} - \boldsymbol{u}_a|^2 \,\mathrm{d}\boldsymbol{v}.$$

The terms Q_{ab} are the energy source terms, representing transfer of energy from species b to species a due to collisions, and are defined by the integral

$$Q_{ab} = \frac{m_a}{2} \int_{\mathbb{R}^d} |\boldsymbol{v}|^2 C_{ab}(f_a, f_b) \,\mathrm{d}\boldsymbol{v}. \tag{2.2.8}$$

We note here that the stress tensor Π_a and the heat flux $\boldsymbol{q}_a$ cannot be determined from the evolved moments

n_a, $n_a\boldsymbol{u}_a$, and e_a. Therefore, to obtain a usable system of equations, we must propose a *closure*, i.e. a way to determine an approximate value for Π_a and $\boldsymbol{q}_a$ from the first five moments. This is often done by treating some asymptotic limit of (2.1.2). One such ansatz is to assume that the distribution function is everywhere Maxwellian, and that interspecies collisions are weak. This yields the **ideal five-moment** equations [94], which capture charge separation and ion inertial effects, but neglect viscous and thermal diffusion as well as collisional resistivity:

$$\begin{aligned} \partial_t n_a + \nabla\cdot(n_a\boldsymbol{u}_a) &= 0 \\ m_a\partial_t(n_a\boldsymbol{u}_a) + \nabla\cdot(m_a n_a\boldsymbol{u}_a\otimes\boldsymbol{u}_a) + \nabla p_a &= n_a q_a(\boldsymbol{E} + \boldsymbol{u}_a\times\boldsymbol{B}) \\ \partial_t e_a + \nabla\cdot((e_a+p_a)\boldsymbol{u}_a) &= n_a q_a\boldsymbol{E}_a\cdot\boldsymbol{u}_a. \end{aligned} \tag{2.2.9}$$

More descriptive fluid models can be obtained by assuming that the distribution function is near a Maxwellian and taking into account interspecies collisions. In this way, non-ideal physics such as viscosity, thermal diffusion and resistivity can be taken into account. The most notable such model is due to Braginskii [11], which assumes strong collisions. In Chapter 3 we develop an asymptotic expansion for (2.1.2) in the absence of collisions which yields expressions for leading-order finite Larmor radius transport effects.

2.3 Maxwell's equations

The equations of plasma dynamics must be coupled to Maxwell's equations in a vacuum, which describe the evolution of the electric and magnetic fields in the rarefied environment of a plasma. Maxwell's equations in a vacuum are written in differential form as follows:

Name	
Ampère's law	$\nabla\times\mathbf{B} = \mu_0(\boldsymbol{j} + \epsilon_0\partial_t\boldsymbol{E})$
Gauss's law	$\nabla\cdot\boldsymbol{E} = \rho_c/\epsilon_0$
Faraday's law	$\nabla\times\boldsymbol{E} = -\partial_t\boldsymbol{B}$
Gauss's law for magnetic fields	$\nabla\cdot\boldsymbol{B} = 0$

The quantities ρ_c and $\boldsymbol{j}$ are the charge density and current density of the plasma, respectively. They are defined as

$$\rho_c = \sum_a q_a n_a, \quad \boldsymbol{j} = \sum_a q_a n_a\boldsymbol{u}_a, \tag{2.3.1}$$

where the sums are taken over all plasma species. The constants μ_0 and ϵ_0 are the vacuum permeability and vacuum permittivity respectively. In SI-MKS units, they have the numerical values

$$\mu_0 = 4\pi \times 10^{-7}\,\mathrm{H/m},$$
$$\epsilon_0 \approx 8.8542 \times 10^{-12}\,\mathrm{F/m},$$

where H denotes the henry and F the farad [88]. The vacuum permeability and permittivity are related by

$$c^2 = \frac{1}{\mu_0 \epsilon_0},$$

where c is the speed of light in a vacuum.

Maxwell's equations are overdetermined, in the sense that if the involutions (Gauss's law and the divergence-of-$\boldsymbol{B}$ constraint) are satisfied at some starting time, then they will be satisfied at all future times by any solution of the pair of time-dependent equations. However, they provide a constraint on valid initial conditions, and they can fail to be satisfied numerically in some cases.

In many applications it is valid to neglect the plasma current and treat $\boldsymbol{B}$ as time-independent given data. In this case, assuming $\nabla \times \boldsymbol{B} = 0$, one need only solve Ampère's law or Gauss's law. This is known as the electrostatic approximation. Roughly speaking, it is valid when the amount of energy available in the plasma to perturb the magnetic field is relatively small; that is, when

$$p \ll \frac{|\boldsymbol{B}|^2}{2\mu_0}.$$

Such plasmas are known as low-beta plasmas, after the dimensionless quantity

$$\beta = \frac{p}{|\boldsymbol{B}|^2/2\mu_0}.$$

In the electrostatic approximation, one typically solves Gauss's law by introducing an electric potential, denoted ϕ, which satisfies

$$\nabla^2 \phi = -\frac{\rho_c}{\epsilon_0},$$

and computes the electric field from $\boldsymbol{E} = -\nabla \phi$.

2.4 Plasma normalizations

A key step in the simulation of physical systems such as plasma is the choice of a normalization. Normalizations relate the numerical quntities in simulation to real physical quantities. The task, in a nutshell, is to formulate a rule for how much “1” of each simulation quantity represents in physical units. We use a slight modification of the flexible plasma normalization described in [83].

The normalization begins by eliminating the vacuum permittivity and permeability. To eliminate ϵ_0, we first introduce a reference number density n_0 and the reference proton plasma frequency

$$\omega_p^2 = \frac{e^2 n_0}{m_p \epsilon_0},$$

where e is the elementary charge and m_p is the proton mass. The vacuum permeability is eliminated via the introduction of a reference Alfvén speed,

$$v_A^2 = \frac{B_0^2}{m_p n_0 \mu_0},$$

which also serves as a reference velocity $v_0 = v_A$. We now introduce a reference length scale L and reference time scale τ, related by $L = v_0 \tau$. The time scale τ may be combined with the plasma frequency and the cyclotron frequency, $\omega_c = \frac{eB_0}{m_p}$, to give a pair of nondimensional quantities $\omega_p \tau$ and $\omega_c \tau$ which characterize the strength of electric and magnetic fields, respectively. Indeed, we set the reference electric field strength to

$$E_0 = \frac{m_p v_0 \omega_p}{e}.$$

We have already set a reference quantity with dimensions of energy, via the magnetic energy density $B_0^2/2\mu_0$. With this in mind, it is natural to fix a reference value for pressure via the plasma β. Setting $\beta = 2$ gives

$$p_0 = n_0 T_0 = \frac{B_0^2}{\mu_0}, \tag{2.4.1}$$

from which we can easily verify $T_0 = m_p v_0^2$. The mass and charge of a species a is eliminated via the reference quantities m_p and e, via

$$m_a = A_a m_p, \quad q_a = Z_a e.$$

Finally, it is useful to introduce a third dimensionless quantity characterizing the strength of collisions, $\nu_p \tau$, where ν_p is the characteristic proton collision frequency:

$$\nu_p = \frac{e^4 n_0 \ln \Lambda}{3\sqrt{2}\pi^{3/2} \epsilon_0^2 m_p^{1/2} T_0^{3/2}}.$$

The characteristic collision frequency enters the normalization via

$$C_{ab} = \nu_p \overline{C}_{ab}.$$

Quantity	Reference value
Mass	$m_p = 1.6726 \times 10^{-27}\,\mathrm{kg}$
Charge	$e = 1.6022 \times 10^{-19}\,\mathrm{C}$
Number density	n_0
Time	τ
Magnetic field	$B_0 = \omega_c m_p / e$
Velocity	$v_0 = \sqrt{B_0^2 / m_p n_0 \mu_0}$
Temperature	$T_0 = \frac{B_0^2}{\mu_0 n_0} = m_p v_0^2$
Length	$L = v_0 \tau$
Electric field	$E_0 = m_p v_0 \omega_p / e$
Collision frequency	$\nu_p = \frac{e^4 n_0 \ln \Lambda}{3\sqrt{2}\pi^{3/2}\epsilon_0^2 m_p^{1/2} T_0^{3/2}}$

Table 2.1: Reference quantities used in the flexible plasma normalization. Exactly three of the values highlighted in red must be supplied, for example as the triple $(n_0, \omega_p\tau, \omega_c\tau)$.

2.4.1 Normalized Vlasov-Boltzmann equation

The normalization proceeds by writing all dimensional quantities ψ as $\psi_0\overline{\psi}$, where ψ_0 is the reference value of that quantity as introduced above, and the barred quantity $\overline{\psi}$ is dimensionless. Derivatives transform according to the chain rule as

$$\partial_t \mapsto \frac{1}{\tau}\partial_{\overline{t}}, \quad \nabla_{\boldsymbol{x}} \mapsto \frac{1}{L}\nabla_{\overline{\boldsymbol{x}}}, \quad \nabla_{\boldsymbol{v}} \mapsto \frac{1}{v_0}\nabla_{\overline{\boldsymbol{v}}}.$$

For kinetic equations, we must also introduce a reference value for the particle distribution function, which is $f_0 = n_0 v_0^{-3}$. Expanding all terms in (2.1.2) in terms of the reference quantities, we obtain

$$\begin{aligned}\frac{f_0}{\tau}\partial_t \overline{f}_a + v_0\overline{\boldsymbol{v}} \cdot \frac{f_0}{L}\nabla_{\boldsymbol{x}}\overline{f}_a + \frac{Z_a e}{A_a m_p}\left(\frac{m_p v_0 \omega_p}{e}\overline{\boldsymbol{E}} + v_0\frac{\omega_c m_p}{e}\overline{\boldsymbol{v}} \times \overline{\boldsymbol{B}}\right) \cdot \frac{1}{v_0}\nabla_{\overline{\boldsymbol{v}}}\overline{f}_a \\ = \nu_p f_0 \sum_b \overline{C}_{ab}(\overline{f}_a, \overline{f}_b).\end{aligned}$$

Multiplying through by τ/f_0 and simplifying, we obtain the following nondimensional Vlasov-Boltzmann equation:

$$\partial_t f_a + \boldsymbol{v} \cdot \nabla_{\boldsymbol{x}} f_a + \frac{Z_a}{A_a}\left[(\omega_p\tau)\boldsymbol{E} + (\omega_c\tau)\boldsymbol{v} \times \boldsymbol{B}\right] \cdot \nabla_{\boldsymbol{v}} f_a = \nu_p\tau \sum C_{ab}(f_a, f_b). \tag{2.4.2}$$

We have removed overlines from dimensionless quantities for clarity. It can be seen that (2.4.2) contains three nondimensional parameters, $\omega_p\tau, \omega_c\tau$, and $\nu_p\tau$, which respectively characterize the strength of electric fields, magnetic fields, and collisions.

The normalized Lenard-Bernstein collision operator $\overline{C}_{ab}$ is computed as follows:

$$
\begin{aligned}
\overline{C}_{ab}(\overline{f}_a, \overline{f}_b) &= \frac{\nu_{ab}}{\nu_p f_0} \nabla_{\boldsymbol{v}} \cdot \left(\frac{T_{ab}}{m_a} \nabla_{\boldsymbol{v}} f_a + (\boldsymbol{v} - \boldsymbol{u}_{ab}) f_a \right) \\
&= n_b \frac{\sqrt{m_a m_b}(m_a + m_b)(q_a q_b)^2 \ln \Lambda}{3\sqrt{2}\pi^{3/2}\epsilon_0^2 m_a (m_a T_b + m_b T_a)^{3/2}} \times \frac{3\sqrt{2}\pi^{3/2}\epsilon_0^2 m_p^{1/2} T_0^{3/2}}{e^4 n_0 \ln \Lambda} \\
&\quad \times \nabla_{\overline{\boldsymbol{v}}} \left(\frac{\overline{T}_{ab}}{A_a} \nabla_{\overline{\boldsymbol{v}}} \overline{f}_a + (\overline{v} - \overline{u}_{ab}) \overline{f}_a \right) \\
&= \overline{n}_b \frac{\sqrt{A_a A_b}(A_a + A_b)(Z_a Z_b)^2}{A_a (A_a \overline{T}_b + A_b \overline{T}_a)^{3/2}} \nabla_{\overline{\boldsymbol{v}}} \left(\frac{\overline{T}_{ab}}{A_a} \nabla_{\overline{\boldsymbol{v}}} \overline{f}_a + (\overline{v} - \overline{u}_{ab}) \overline{f}_a \right).
\end{aligned}
$$

Rewriting without overlines,

$$
C_{ab}(f_a, f_b) = \nu_{ab} \nabla_{\boldsymbol{v}} \cdot \left(\frac{T_{ab}}{m_a} \nabla_{\boldsymbol{v}} f_a + (\boldsymbol{v} - \boldsymbol{u}_{ab}) f_a \right)
$$

where

$$
\nu_{ab} = n_b \frac{\sqrt{A_a A_b}(A_a + A_b)(Z_a Z_b)^2}{A_a (A_a T_b + A_b T_a)^{3/2}}. \tag{2.4.3}
$$

2.4.2 Normalized Maxwell's equations

Maxwell's equations may be normalized in the same way as the Vlasov-Boltzmann equation. For Ampère's law, we introduce the reference current density j_0 as

$$
j_0 = e v_0 n_0.
$$

We also eliminate ϵ_0 and μ_0 by using the expression for reference Alfvén velocity and plasma frequency, which give

$$
\epsilon_0 = \frac{e^2 n_0}{\omega_p^2 m_p}, \quad \mu_0 = \frac{B_0^2}{p_0} = \frac{B_0^2}{n_0 m_p v_0^2}.
$$

Substituting the above into Ampère's law, we get

$$
\begin{aligned}
\frac{B_0}{L} \nabla_{\overline{\boldsymbol{x}}} \times \overline{\boldsymbol{B}} &= \frac{B_0^2}{n_0 m_p v_0^2} \left(e v_0 n_0 \overline{\boldsymbol{j}} + \frac{e^2 n_0}{\omega_p^2 m_p} \frac{1}{\tau} \frac{m_p v_0 \omega_p}{e} \partial_{\overline{t}} \overline{\boldsymbol{E}} \right) \\
&= \frac{B_0^2}{n_0 m_p v_0} e v_0 n_0 \left(\overline{\boldsymbol{j}} + \frac{1}{\omega_p \tau} \partial_{\overline{t}} \overline{\boldsymbol{E}} \right) \\
&= \frac{B_0 \omega_c}{v_0} \left(\overline{\boldsymbol{j}} + \frac{1}{\omega_p \tau} \partial_{\overline{t}} \overline{\boldsymbol{E}} \right).
\end{aligned}
$$

Cancelling the factor of B_0/L and removing overlines gives the normalized Ampère's law

$$\nabla \times \boldsymbol{B} = \omega_c \tau \left(\boldsymbol{j} + \frac{1}{\omega_p \tau} \partial_t \boldsymbol{E} \right). \tag{2.4.4}$$

Gauss's law is normalized as follows:

$$\frac{1}{L} \nabla_{\overline{\boldsymbol{x}}} \cdot E_0 \overline{\boldsymbol{E}} = e n_0 \frac{\overline{\rho}_c}{\epsilon_0}.$$

Using the expression for ϵ_0 in terms of plasma frequency, we obtain

$$\begin{aligned} \frac{1}{L} \nabla_{\overline{\boldsymbol{x}}} \cdot \frac{m_p \omega_p v_0}{e} \overline{\boldsymbol{E}} &= \frac{\omega_p^2 m_p}{e^2 n_0} e n_0 \overline{\rho}_c \\ \nabla_{\overline{\boldsymbol{x}}} \cdot \overline{\boldsymbol{E}} &= (\omega_p \tau) \overline{\rho}_c, \end{aligned}$$

or, removing overlines for clarity,

$$\nabla \cdot \boldsymbol{E} = (\omega_p \tau) \rho_c. \tag{2.4.5}$$

For Faraday's law, we write

$$\frac{1}{L} \nabla_{\overline{\boldsymbol{x}}} \times \frac{m_p v_0 \omega_p}{e} \overline{\boldsymbol{E}} = -\frac{1}{\tau} B_0 \partial_{\overline{t}} \overline{\boldsymbol{B}}.$$

Multiplying through by $L\tau e / v_0$, we get

$$\omega_p \tau \nabla \times \boldsymbol{E} = -\omega_c \tau \partial_t \boldsymbol{B}. \tag{2.4.6}$$

Finally, the normalized form of the divergence law for the magnetic field is unchanged:

$$\nabla \cdot \boldsymbol{B} = 0. \tag{2.4.7}$$

Together, (2.4.4), (2.4.5), (2.4.6), and (2.4.7) form the normalized Maxwell's equations.

2.4.3 Use of the flexible normalization

It is worth dwelling for a moment on how to use the normalization described in this section, not least because the question confused this author for quite some time. The reference quantities listed in Table 2.1 are uniquely determined by the true physical values of the constants of nature ($\epsilon_0, \mu_0, m_p, m_e$ and e), *up to three degrees of freedom.* It is up to the user of the normalization to supply three quantities which, in combination with the constants of nature, determine dimensional values for each of the reference values.

The combination of three such values is called a *normalization triplet.* An example of a viable normalization triplet is $(n_0, \omega_p\tau, \omega_c\tau)$. From n_0, the reference proton plasma frequency is uniquely determined by

$$\omega_p^2 = \frac{n_0 e^2}{\epsilon_0 m_p},$$

from which one can determine τ and thence ω_c. From ω_c, the reference magnetic field is given by $B_0 = \omega_c m_p/e$, which sets a reference energy and Alfvén speed, and so on.

For completeness, we will show how to use a different normalization triplet which is useful for controlling the relationship between magnetization and collisionality, namely $(n_0, \nu_p\tau, \omega_c\tau)$. The normalization procedure with this triplet begins as follows:

$$\begin{aligned}
\left(B_0^2\tau^2\right) &= (\omega_c\tau)^2 \frac{m_p^2}{e^2} \\
T_0\tau^2 &= n_0^{-1}\frac{B_0^2\tau^2}{\mu_0} \\
\frac{\tau}{T_0^{3/2}} &= \nu_p\tau \frac{3\sqrt{2}\pi^{3/2}\epsilon_0^2 m_p^{1/2}}{n_0 e^4 \ln\Lambda} \\
T_0^{-4} &= \left(\frac{\tau}{T_0^{3/2}}\right)^2 (T_0\tau^2)^{-1}.
\end{aligned}$$

Once T_0 is determined, τ follows easily from $\nu_p\tau$ and the remainder of the normalization calculations are straightforward.

$$\begin{aligned}
n_0\tau^2 &= (\omega_p\tau)^2 \frac{\epsilon_0 m_p}{e^2} \\
\left(B_0^2\tau^2\right) &= (\omega_c\tau)^2 \frac{m_p^2}{e^2} \\
T_0 &= \frac{B_0^2\tau^2}{\mu_0 n_0\tau^2} \\
\tau^{-1} &= \frac{\nu_p\tau}{n_0\tau^2} T_0^{3/2} \frac{3\sqrt{2}\pi^{3/2}\epsilon_0^2 m_p^{1/2}}{e^4 \ln\Lambda}
\end{aligned}$$

The plasma normalization serves to relate the dimensionless values which we manipulate in simulation to physical dimensions. This relation is in one-to-one correspondence with the dimensionless form of the plasma equations such as (2.4.2). It is correct to think of the normalization triplet as determining some aspects of the plasma “regime” under simulation. The value of $\nu_p\tau$ that results from the reference calculations certainly determines whether the plasma is in a collisional or a collisionless regime. However, the normalization triplet does not determine all aspects of the plasma “regime”. For example the plasma beta is independent of the normalization triplet. A simulated plasma with $\overline{T} = 1$, $\overline{B} = 1$ has $\beta = 2$, while $\overline{T} = 10^{-4}$, $\overline{B} = 1$ indicates a β so small that the electrostatic approximation is likely valid.

One way to view the idea of normalization is as fixing a particular nondimensional Vlasov-Boltzmann or fluid equation. Different equations certainly have different solutions, but there is just as much variety in the set of solutions to a single nonlinear equation depending on initial values.

Chapter 3

Non-ideal perpendicular transport of low-beta plasmas

As we have seen, the derivation of moment-based fluid models from kinetic models leads to the infamous "closure problem", which is the appearance of a flux of moment $N+1$ in the equation for moment N. For a usable moment method, this closure problem must be solved through the introduction of some suitable ansatz for f, typically physically motivated. The moment methods we will consider here use a Hilbert expansion ansatz,

$$f = f^0 + \epsilon f^1 + \epsilon^2 f^2 + \cdots ,$$

where ϵ is a small parameter providing scale separation between dominant and perturbative physics. One such ansatz is provided by the assumption of strong collisions, namely that f is near a Maxwellian: $f^0 = \mathcal{M}$. Once an ansatz for f is chosen, careful asymptotic analysis, wherein we retain terms of order ϵ and discard higher-order terms, can provide expressions for the non-ideal moments in the momentum and energy equations. This is done through the Chapman-Enskog procedure, which involves solving for f^1 in terms of the moments of f^0.

In the classic paper "Transport Processes in a Plasma", Braginskii laid out a Chapman-Enskog type procedure for a simple plasma consisting of one ion and one electron species [11]. The key insight in Braginskii's work, beyond previous attempts at Chapman-Enskog analyses for multispecies gases, was that the large ion to electron mass ratio provides a significant scale separation, characterized by a small parameter $\epsilon = \sqrt{m_e/m_i}$. Taking advantage of this smallness, which implies a corresponding scale separation between the electron collisional timescales and the ion collisional timescales, as well as a scale separation between the electron *thermal* velocity and drift velocities, leads to a system of transport closures that supports separate species velocities and temperatures.

Degond [25] later put the Braginskii transport system on a somewhat firmer mathematical footing by precisely retaining exactly those terms of order ϵ and performing the asymptotic analysis at the level of the kinetic equation. The resulting system of equations, *at zeroth order*, is a typical five-moment system for the ion species, and a particle and energy diffusion equation for the electrons. Degond did not expand the order ϵ terms for the ion species.

A shortcoming of the above-mentioned works is that they are based on a strongly collisional assumption, whereas many plasmas of interest are effectively collisionless: the particles undergo negligible scattering over the dynamic timescales of interest. However, magnetization provides a strong regularizing effect for the perpendicular components of the particle distribution function. The tokamak modeling community has long realized the possibility of deriving reduced models for collisionless, magnetized plasmas. One approach, gyrokinetics, involves passing to the guiding center frame of reference via a coordinate transformation [12]. The resulting equation system is kinetic in the parallel and perpendicular particle energies. Gyrofluid models [28] apply a moment closure procedure to the parallel velocity dependence based on kinetic considerations such as the linear electrostatic response.

Other authors have arrived at fluid theories for collisionless, magnetized plasmas. The Chew, Goldberger and Low (CGL) [17] system contains the leading-order behavior of a magnetized plasma's fluid moments, but neglects non-ideal (that is, non-Maxwellian) moments. Macmahon [80] derived the finite-Larmor radius corrections to the CGL equations. In the tokamak modeling community, a similar ordering has been used to derive reduced magnetohydrodynamic equations [57, 61]. Simakov and Catto used a Maxwellian assumption on the leading-order distribution function to derive perpendicular transport [95]. Finally, a comprehensive study of the collisionless and weakly collisional asymptotic expansions of the Vlasov equation was made by Ramos [86, 87], who clarified the relationship between the drift, or slow dynamics, ordering considered here, and the fast dynamics ordering of Braginskii.

With this chapter we present another derivation of the perpendicular fluid behavior of a highly magnetized plasma. In contrast to prior collisionless [80, 86] and weakly collisional [87, 95] expansions, we take an explicitly kinetic approach in the parallel direction by integrating only over perpendicular velocity $\boldsymbol{v}_\perp$. We call the resulting equations "semi-fluid" equations. The semi-fluid equations are kinetic in $v_\parallel$, and evolve recognizable perpendicular moments such as density and total perpendicular energy. Our moment system omits moments like the parallel pressure, as well as more complicated constructions such as the parallel flux of perpendicular energy. Considering only perpendicular moments simplifies the resulting equations, which are kinetic in parallel velocity. The benefit is that the parallel coordinate may either be modeled "intact", or tackled with a different reduced model, such as a gyrofluid system or a dynamical low-rank (DLR) approximation. We will discuss the latter starting in Chapter 6.

The idea of treating the parallel direction kinetically and using a moment expansion for the perpendicular direction is also not without precedent. Classical drift kinetics [56] can be viewed as such a derivation in an ordering where the leading-order perpendicular velocity comprises drifts. The parallel-kinetic perpendicular-moment (PKPM) method developed in [69] is another example, with a focus on solving the perpendicular moment equations using a truncated spectral expansion. Finally, Ref. [110] makes use of a similar idea to calculate a moment closure for the collisionless Vlasov equation in the fast dynamics ordering of Braginskii.

Compared to [110], our expansion is performed in the drift ordering. We therefore obtain a different expression for gyroviscous stress than fast dynamics ordering expansions. The significance of this discrepancy is investigated in detail via numerical experiments in Chapter 5. Furthermore, we are careful to retain collisional moments in an abstract form, thereby indicating how collisional terms may be included in a more complete calculation.

On the benefits of a momentum equation: diffusive and hyperbolic expansions

The most significant difference between the Braginskii expansion for electrons, and Degond's clarification of Braginskii's work [25], is that Braginskii retains an electron momentum equation. This is accomplished by declining to take advantage of all of the smallness available in the asymptotic expansion. That is, at certain points, Braginskii makes use of the fact that the electron drift velocity is order ϵ compared to their thermal velocity. This would allow him to use a gyrotropic ansatz for f_e^0, and locate all electron velocity in f_e^1, thereby obtaining a similar result to Degond. However, he opts to choose a solution for f_e^0 with nonzero velocity, which results in an evolution equation for the electron momentum rather than a generalized Ohm's law.

The significance of this decision for the electron species is not as important as it is for the ion species. Whereas electron inertia is often totally negligible, such that a Boltzmann electron approximation is often used, to neglect ion inertia is to neglect all plasma inertia completely. Indeed, ion inertial effects such as the polarization drift have been found to play an important role in plasma phenomena such as the Kelvin-Helmholtz instability, even in low-beta plasmas [108]. We therefore want an asymptotic expansion that retains at least an ion momentum equation and ion inertial effects. This is achieved by what we call the "hyperbolic expansion", which is a five-moment semi-fluid system of equations with a closure term similar to that derived by Braginskii [11] for the diamagnetic heat flux. On the other hand, as mentioned above, the drift ordering chosen here leads to a different closure for the stress tensor than obtained by Braginskii.

More generally, keeping an eye towards the actual implementation of the transport closure scheme in a numerical algorithm, we are skeptical of differential operators of degree higher than two. A key benefit of a momentum equation is that it allows the inclusion of viscosity without requiring a third-order differential operator on the right-hand side of a numerical scheme. In contrast, a system that eagerly substitutes the leading-order drift velocity for $\boldsymbol{u}$ will require third-order spatial derivatives to capture viscous effects.

This chapter presents two different expansion procedures, one with a momentum equation and one without. We refer to the expansion without a momentum equation, which consists of a continuity and energy equation, as the “diffusive expansion,” since the resulting semi-fluid system can be viewed as a pair of diffusion equations for particles and energy. This expansion results from eagerly solving for $\boldsymbol{u}_\perp$ at each order and eliminating its appearances in the continuity and energy equations. The diffusive expansion is described in Section 3.3.

The other expansion we term the “hyperbolic expansion,” because it takes the form of a hyperbolic system of equations for mass, momentum and energy, similar to the Euler equations. It is obtained by rearranging the order ϵ^0 kinetic equation to locate all momentum in the order ϵ^0 kinetic solution, just as we locate all mass and energy in the order ϵ^{-1} kinetic solution. The hyperbolic expansion is described in Section 3.4.

Organization of the chapter

The remainder of this chapter is organized as follows. Section 3.1 rigorously derives a dimensionless Vlasov-Boltzmann equation in the low-beta regime with appropriate scaling. In Section 3.2, we introduce the perpendicular semi-fluid equation system and the asymptotic analysis framework we will use, including Hilbert expansions and scaling assumptions. Section 3.3 carries the diffusive expansion procedure through to first order in ϵ, and in particular derives an order ϵ closure for the diamagnetic heat flux. In Section 3.4, we modify the expansion procedure to retain a momentum equation and derive an order ϵ^2 closure for the stress tensor through the hyperbolic expansion. Section 3.5 summarizes the two procedures and the resulting systems of equations. Finally, in Section 3.6 we discuss the implications of the analysis in this chapter, and its relation to other reduced models for magnetized plasmas.

3.1 Scaling assumptions

In this section we make precise our normalization and scaling assumptions. Our normalization is based on the flexible plasma normalization described in Section 2.4, but with the introduction of a small parameter ϵ at an appropriate point. Beginning with the dimensional Vlasov equation for species a, (2.1.2), we normalize the species charge and mass to the proton charge e and mass m_p:

$$m_a = A_a m_p, \quad q_a = Z_a e.$$

The reference proton plasma and cyclotron frequencies are given by

$$\omega_p^2 = \frac{e^2 n_0}{m_p \epsilon_0}, \quad \omega_c = \frac{\epsilon B_0 e}{m_p},$$

where n_0 is a reference number density. Note the introduction of a small parameter ϵ in the definition of the reference cyclotron frequency. The plasma frequency eliminates ϵ_0; to eliminate μ_0 we introduce the reference Alfvén velocity

$$v_A^2 = \frac{B_0^2}{m_p n_0 \mu_0}.$$

The reference velocity is set to $v_0 = \epsilon v_A$. We introduce characteristic length and time scales via

$$\boldsymbol{x} = L\overline{\boldsymbol{x}}, \quad t = \tau\overline{t},$$

where τ is a reference timescale and $L = v_0\tau$. Finally, reference phase space densities and collision operators are introduced via

$$f_a = f_0\overline{f}_a, \quad C(f_a) = \nu_p f_0 \overline{C}(\overline{f}_a).$$

The reference quantities are used to nondimensionalize (2.1.3) by substituting expressions such as $\boldsymbol{v} = v_0\overline{\boldsymbol{v}}$, where the notational convention is that overlined quantities are of order unity. Doing so gives

$$\frac{f_0}{\tau}\partial_{\overline{t}}\overline{f}_a + \frac{f_0 v_0}{L}\overline{\boldsymbol{v}}\cdot\nabla\overline{f}_a + \frac{Z_a e}{A_a m_p}\left[\frac{m_p v_0 \omega_p}{e}\overline{\boldsymbol{E}} + v_0\frac{\epsilon^{-1} m_p \omega_c}{e}(\overline{\boldsymbol{v}}\times\overline{\boldsymbol{B}})\right]\cdot\frac{f_0}{v_0}\nabla_{\overline{\boldsymbol{v}}}\overline{f}_a = \nu_p f_0\overline{C}(\overline{f}_a).$$

Multiplying through by τ/f_0 gives the nondimensional Vlasov equation in strongly magnetized scaling,

$$\partial_t f_a + \boldsymbol{v}\cdot\nabla f_a + \frac{Z_a}{A_a}\left[\omega_p\tau\boldsymbol{E} + \epsilon^{-1}\omega_c\tau(\boldsymbol{v}\times\boldsymbol{B})\right]\cdot\nabla_{\boldsymbol{v}} f_a = \nu_p\tau C(f_a). \tag{3.1.1}$$

Overlines are omitted in (3.1.1) and in all subsequent expressions for clarity.

Equation (3.1.1) is expressed in the flexible normalization form, which is characterized by three dimensionless parameters $\omega_p\tau$, $\omega_c\tau$, and $\nu_p\tau$. These characterize the strength of electrostatic forces, magnetic forces, and collisions, respectively. Equation (3.1.1) is additionally equipped with a formal small parameter ϵ around which we will perform asymptotic expansion in the following section. Note that we have separated the contribution of $\omega_c\tau$, which we will constrain to be of order unity, from ϵ^{-1}. Because our focus is on the scaling of the magnetic field term, in what follows we set $\omega_p\tau = \nu_p\tau = 1$ for clarity.

Before proceeding, we first write some important plasma parameters in terms of ϵ. The reference plasma temperature is defined as $T_0 = m_p v_0^2$, and the reference pressure $p_0 = n_0 T_0$. Thus the plasma beta scales as ϵ^2:

$$\beta_0 = \frac{p_0}{B_0^2/2\mu_0} = \frac{n_0 m_p \epsilon^2 v_A^2}{B_0^2/2\mu_0} = 2\epsilon^2.$$

The nondimensional proton Larmor radius is

$$\frac{r_{Li}}{L} = \frac{v_0 m_p}{LeB_0} = \frac{\epsilon v_0}{L\omega_c} = \frac{\epsilon}{\omega_c \tau}.$$

3.2 Parallel-perpendicular hybrid model with closure

We will derive an asymptotic expansion resulting from (3.1.1) in the context of a simplified magnetic geometry. Let $\boldsymbol{B}$ be spatially uniform and constant in time, so that $\boldsymbol{B} = B\hat{\boldsymbol{b}}$ for a constant unit vector $\hat{\boldsymbol{b}}$. The spatial and velocity coordinates can be split into parallel and perpendicular components relative to $\hat{\boldsymbol{b}}$, like so:

$$x_\parallel = \boldsymbol{x} \cdot \hat{\boldsymbol{b}}, \quad \boldsymbol{x}_\perp = -(\boldsymbol{x} \times \hat{\boldsymbol{b}}) \times \hat{\boldsymbol{b}}$$

$$v_\parallel = \boldsymbol{v} \cdot \hat{\boldsymbol{b}}, \quad \boldsymbol{v}_\perp = -(\boldsymbol{v} \times \hat{\boldsymbol{b}}) \times \hat{\boldsymbol{b}}.$$

We then can write $\boldsymbol{v} = (v_\parallel, \boldsymbol{v}_\perp)$ and $\boldsymbol{x} = (x_\parallel, \boldsymbol{x}_\perp)$. We also introduce the parallel and perpendicular gradient operators

$$\partial_\parallel = \frac{\partial}{\partial x_\parallel}, \quad \nabla_\perp = \nabla_{\boldsymbol{x}_\perp}.$$

Our method begins by deriving a set of reduced equations, what we call semi-fluid equations, by taking moments of (3.1.1) in $\boldsymbol{v}_\perp$. In this way we obtain semi-fluid analogues to the usual $5N$-moment models, involving equations for density, momentum and energy. By semi-fluid, we mean that the evolved quantities themselves depend on $v_\parallel$ in addition to space and time. This is therefore a model reduction from 3 to 1 velocity variables.

The result of the moment-taking procedure is an exact hierarchy of moment equations posed over $\boldsymbol{x}, t$, and unconventionally, $v_\parallel$. This moment hierarchy presents the typical moment-closure problem: cutting off the hierarchy after the energy equation leaves undetermined quantities, namely the perpendicular heat flux and full perpendicular pressure tensor. The closure problem is addressed by introducing a Hilbert expansion for f_a centered around a gyrotropic Maxwellian. Higher-order corrections are obtained by inverting the leading-order operator, which in the case of (3.1.1) is simply the cyclotron rotation operator.

The Chapman-Enskog hierarchy resulting from this procedure is distinct from the hierarchy of asymptotic corrections in Braginskii [11] in two ways. To begin with, we will show that the leading-order solution f^0 cannot support a nonzero drift velocity, since it must be gyrotropic. The leading-order drift velocity therefore appears at order ϵ, rather than as a parameter of a leading-order Maxwellian distribution. This choice of ordering has been called the "drift" ordering, or the "slow-dynamics" ordering, in contrast to the fast-dynamics ordering where the fluid velocity is taken to be large [86].

The second departure from Braginskii makes the required calculations significantly easier: the collision operator appears explicitly in each subsequent correction equation, and does not need to be inverted at each step. As we will see, this makes it far easier to investigate the collisional contributions to transport for different collision operators.

Semi-fluid equations

We define the following family of perpendicular moments. In these and most subsequent equations, we omit the species subscript a, except when called for to distinguish between species as in the collisional terms.

$$n_\perp(\boldsymbol{x}, v_\parallel, t) = \int f \, \mathrm{d}\boldsymbol{v}_\perp, \tag{3.2.1}$$

$$\boldsymbol{u}_\perp(\boldsymbol{x}, v_\parallel, t) = \frac{1}{n_\perp} \int \boldsymbol{v}_\perp f \, \mathrm{d}\boldsymbol{v}_\perp \tag{3.2.2}$$

$$\mathbb{P}_\perp(\boldsymbol{x}, v_\parallel, t) = A \int (\boldsymbol{v}_\perp - \boldsymbol{u}_\perp) \otimes (\boldsymbol{v}_\perp - \boldsymbol{u}_\perp) f \, \mathrm{d}\boldsymbol{v}_\perp \tag{3.2.3}$$

$$e_\perp(\boldsymbol{x}, v_\parallel, t) = \frac{A}{2} \int |\boldsymbol{v}_\perp|^2 f \, \mathrm{d}\boldsymbol{v}_\perp \tag{3.2.4}$$

$$\boldsymbol{q}_\perp(\boldsymbol{x}, v_\parallel, t) = \frac{A}{2} \int (\boldsymbol{v}_\perp - \boldsymbol{u}_\perp) |\boldsymbol{v}_\perp - \boldsymbol{u}_\perp|^2 \, \mathrm{d}\boldsymbol{v}_\perp \tag{3.2.5}$$

$$\mathcal{N}(\boldsymbol{x}, v_\parallel, t) = \int C(f) \, \mathrm{d}\boldsymbol{v}_\perp \tag{3.2.6}$$

$$\mathcal{S}(\boldsymbol{x}, v_\parallel, t) = A \int \boldsymbol{v}_\perp C(f) \, \mathrm{d}\boldsymbol{v}_\perp \tag{3.2.7}$$

$$\mathcal{Q}(\boldsymbol{x}, v_\parallel, t) = \frac{A}{2} \int |\boldsymbol{v}_\perp|^2 C(f) \, \mathrm{d}\boldsymbol{v}_\perp. \tag{3.2.8}$$

By taking the zeroth moment of (3.1.1) in $\boldsymbol{v}_\perp$, we get the semi-fluid continuity equation:

$$\frac{\mathrm{d}}{\mathrm{d}t} n_\perp + \nabla_\perp \cdot (n_\perp \boldsymbol{u}_\perp) = \mathcal{N}, \tag{3.2.9}$$

where the total derivative operator $\frac{\mathrm{d}}{\mathrm{d}t}$ is defined

$$\frac{\mathrm{d}}{\mathrm{d}t} = \partial_t + \mathcal{D}_\parallel,$$

with

$$\mathcal{D}_\parallel = v_\parallel \partial_\parallel + \frac{Z}{A} E_\parallel \partial_{v_\parallel}.$$

Equation (3.2.9) describes the evolution of the density of particles with a given parallel velocity in space.

Name	Moment	Equation
Continuity	$\int \cdot \, \mathrm{d}\boldsymbol{v}_\perp$	$\frac{\mathrm{d}}{\mathrm{d}t} n_\perp + \nabla_\perp \cdot (n_\perp \boldsymbol{u}_\perp) = \mathcal{N}$
Momentum	$A \int \cdot \boldsymbol{v}_\perp \, \mathrm{d}\boldsymbol{v}_\perp$	$A \frac{\mathrm{d}}{\mathrm{d}t}(n_\perp \boldsymbol{u}_\perp) + \nabla_\perp \cdot (A n_\perp \boldsymbol{u}_\perp \otimes \boldsymbol{u}_\perp + \mathbb{P}_\perp) = n_\perp Z(\boldsymbol{E}_\perp + \epsilon^{-1}\omega_c \tau \boldsymbol{u}_\perp \times \boldsymbol{B}) + \mathcal{S}$
Energy	$\frac{A}{2} \int \cdot \lvert \boldsymbol{v}_\perp \rvert^2 \, \mathrm{d}\boldsymbol{v}_\perp$	$\frac{\mathrm{d}}{\mathrm{d}t} e_\perp + \nabla_\perp \cdot ((e_\perp \mathbb{I} + \mathbb{P}_\perp) \cdot \boldsymbol{u}_\perp + \boldsymbol{q}_\perp) = n_\perp Z \boldsymbol{E}_\perp \cdot \boldsymbol{u}_\perp + \mathcal{Q}$

Table 3.1: Summary of semi-fluid equations

It resembles a typical fluid continuity equation in the perpendicular direction, but is governed by a Vlasov operator in the parallel direction. It is important to note the presence of a source term on the right-hand side, $\mathcal{N}$, which is the zeroth $\boldsymbol{v}_\perp$ moment of the collision operator. Whereas elastic collisions cannot create or destroy particles, and therefore the zeroth moment over all velocity space of the collision operator vanishes, we must account for those particles which arrive at a particular $v_\parallel$ as the result of collisions. The statement that collisions conserve particle number is

$$\int \mathcal{N} \, \mathrm{d}v_\parallel = 0.$$

The collisional particle source term $\mathcal{N}$ is an unclosed moment quantity which must be determined to obtain a self-consistent semi-fluid system.

The perpendicular momentum equation is obtained by taking the first perpendicular moment of (3.1.1), which gives

$$A \frac{\mathrm{d}}{\mathrm{d}t}(n_\perp \boldsymbol{u}_\perp) + \nabla_\perp \cdot (A n_\perp \boldsymbol{u}_\perp \otimes \boldsymbol{u}_\perp + \mathbb{P}_\perp) = n_\perp Z(\boldsymbol{E}_\perp + \epsilon^{-1}\omega_c \tau \boldsymbol{u}_\perp \times \boldsymbol{B}) + \mathcal{S} \tag{3.2.10}$$

Once again, the flux of perpendicular momentum in the parallel direction is governed by a simple free streaming operator. The momentum equation contains another two unclosed moments, which are the full perpendicular pressure tensor $\mathbb{P}_\perp$ and a collisional term $\mathcal{S}$ which deserves some explanation. Being a moment of the collision term, $\mathcal{S}$ represents a collisional transfer of momentum. However, in addition to the usual frictional force term, which describes a transfer of momentum between two species, $\mathcal{S}$ contains a within-species collisional flux of perpendicular momentum between locations in $v_\parallel$ space.

Finally, the perpendicular energy equation is

$$\frac{\mathrm{d}}{\mathrm{d}t} e_\perp + \nabla_\perp \cdot ((e_\perp \mathbb{I} + \mathbb{P}_\perp) \cdot \boldsymbol{u}_\perp + \boldsymbol{q}_\perp) = n_\perp Z \boldsymbol{E}_\perp \cdot \boldsymbol{u}_\perp + \mathcal{Q}. \tag{3.2.11}$$

Here we see the remaining pair of unclosed quantities: the perpendicular flux of (perpendicular) heat, $\boldsymbol{q}_\perp$, and the collisional transfer of perpendicular energy $\mathcal{Q}$. As with the momentum transfer term $\mathcal{S}$, $\mathcal{Q}$ combines inter-species energy transfer and within-species transfer of energy between locations in $v_\parallel$ space.

Hilbert expansions

The next step is to introduce an asymptotic expansion which will allow us to close the moment equations summarized in Table 3.1. This takes the form of a Hilbert expansion, in which unknowns are expanded in powers of the small parameter ϵ. To analyze the kinetic solution hierarchy we need a Hilbert expansion for f, which is

$$f = f^0 + \epsilon f^1 + \epsilon^2 f^2 + \cdots . \tag{3.2.12}$$

Each of the macroscopic parameters is also endowed with a Hilbert expansion, for which we use the same notation. For example,

$$\boldsymbol{u}_\perp = \boldsymbol{u}_\perp^0 + \epsilon \boldsymbol{u}_\perp^1 + \epsilon^2 \boldsymbol{u}_\perp^2 + \cdots . \tag{3.2.13}$$

In order to leave the treatment of collisions until later, it will also be convenient to supply a Hilbert expansion for the collisional terms, such as

$$C(f) = C^0(f) + \epsilon C^1(f) + \epsilon^2 C^2(f) + \cdots .$$

Finally, we introduce the concept of slow time and medium time scales. Letting $t = t^0 + \epsilon^{-1} t^1$, we have

$$\partial_t = \partial_{t^0} + \epsilon \partial_{t^1}.$$

We use the term "medium" timescale to contrast with the fastest timescale, which is the cyclotron frequency timescale. In the course of the formal asymptotic expansion, we will find that the medium timescale denoted by t^0 contains dynamics in the parallel direction, while the slow timescale t^1 contains dynamics of agyrotropic perpendicular motion. We also define the total derivative at the medium timescale t^0 as

$$\frac{\mathrm{d}}{\mathrm{d}t^0} = \partial_{t^0} + \mathcal{D}_\parallel .$$

3.2.1 Order-ϵ^{-1} equation

To begin with, we substitute (3.2.12) into (3.1.1) to obtain the order-ϵ^{-1} kinetic equation,

$$\frac{Z}{A} \omega_c \tau \boldsymbol{v}_\perp \times \boldsymbol{B} \cdot \nabla_{\boldsymbol{v}} f^0 = 0. \tag{3.2.14}$$

This has general solutions which are gyrotropic in $\boldsymbol{v}_\perp$, that is, functions of $v_\parallel$ and $|\boldsymbol{v}_\perp|^2$. However, in this work we assume that the leading order solution is a Maxwellian,

$$f^0 = \mathcal{M} = \frac{A{n_\perp}^0(\boldsymbol{x}, v_\parallel, t)}{2\pi T_\perp^0(\boldsymbol{x}, v_\parallel, t)} \exp\left(-\frac{A|\boldsymbol{v}_\perp|^2}{2T_\perp^0(\boldsymbol{x}, v_\parallel, t)}\right). \tag{3.2.15}$$

It is important to recognize that this is a modeling assumption. The nature of the governing equation (3.1.1) is that its leading-order solutions are gyrotropic, not necessarily Maxwellian. However, it may be justified to assume that f^0 is a gyrotropic Maxwellian in certain cases, particularly when flow velocities are not too large. As we will see, in the presence of moderate collisions, this assumption may be replaced by an assumption that the medium-time evolution is negligible, $\partial_{t^0} = 0$. In either case, the assumption of a Maxwellian leading-order solution may or may not match any particular physical situation.

We assume that the density of $\mathcal{M}$ is equal to $n_\perp$. In other words, $n_\perp^1 = n_\perp^2 = \cdots = 0$. This is accomplished simply by choosing a solution of (3.2.14) whose density is equal to $n_\perp$. A similar assumption is required for the temperature parameter. In the fast dynamics ordering, such as Braginskii's expansion, it is natural to assume $T_\perp^0 = T_\perp$, i.e. that the Maxwellian's temperature is equal to the total temperature $T_\perp$. Here, we make the choice that the Maxwellian's *energy* is equal to the total energy of f: $p_\perp^0 = n_\perp T_\perp^0 = e_\perp$. The result is that higher-order corrections to the temperature and scalar pressure will appear in subsequent equations.

We highlight a major departure from standard asymptotic expansions for the Vlasov equation, which is that the leading-order drift velocity is zero:

$$\boldsymbol{u}_\perp^0 = \int \boldsymbol{v}_\perp f^0 \,\mathrm{d}\boldsymbol{v}_\perp = 0.$$

This same fact was derived by Degond [25] for the electron species, which in that work are treated as subject to collisional effects at the same order as the Lorentz force. In Braginskii's expansion, the electron drift velocity is treated as small (relative to the electron thermal velocity) to eliminate certain other terms, but the leading-order distribution f_e is given a nonzero drift velocity.

Some important properties of the leading order solution (3.2.15) are

$$\int \mathcal{M}\,\mathrm{d}\boldsymbol{v}_\perp = n_\perp, \quad \frac{A}{2}\int |\boldsymbol{v}_\perp|^2 \mathcal{M}\,\mathrm{d}\boldsymbol{v}_\perp = n_\perp T_\perp^0 = p_\perp^0 = e_\perp \tag{3.2.16}$$

and

$$\begin{aligned}\nabla_\perp \mathcal{M} &= \left(\frac{\nabla_\perp n_\perp}{n_\perp} - \frac{\nabla_\perp T_\perp^0}{T_\perp^0} + \frac{A|\boldsymbol{v}_\perp|^2 \nabla_\perp T_\perp^0}{2(T_\perp^0)^2}\right)\mathcal{M} && (3.2.17)\\ &\triangleq \boldsymbol{p}\mathcal{M} + \frac{A|\boldsymbol{v}_\perp|^2}{T_\perp^0}\boldsymbol{r}\mathcal{M}, && (3.2.18)\end{aligned}$$

where we have defined the vectors

$$\boldsymbol{p} = \frac{\nabla n_\perp}{n_\perp} - \frac{\nabla T_\perp^0}{T_\perp^0}, \quad \boldsymbol{r} = \frac{\nabla T_\perp^0}{2T_\perp^0}.$$

3.3 Diffusive expansion

In this section, we carry out what we call the diffusive expansion, which omits a momentum equation and uses an equation of state to solve for the leading-order drift velocity. We will obtain a closure for the perpendicular heat flux $\boldsymbol{q}_\perp$ to leading order.

3.3.1 Semi-fluid equations to order ϵ

Having defined the zeroth-order quantities, we now write down a system of semi-fluid equations to order ϵ, some of whose terms will require closure. The character of our equation system is substantially different than the classic $5N$-moment fluid system. Because the zeroth-order solution has no drift velocity, the semi-fluid momentum equation will play the role of an equation of state.

Recall first that $\boldsymbol{u}_\perp^0 = 0$. By discarding terms of order ϵ in the semi-fluid momentum equation (3.2.10), we obtain an equation of state for $\boldsymbol{u}_\perp^1$:

$$\nabla_\perp p_\perp^0 = n_\perp Z(\boldsymbol{E}_\perp + \omega_c \tau \boldsymbol{u}_\perp^1 \times \boldsymbol{B}). \tag{3.3.1}$$

We have neglected $\mathcal{S}^0$: since the leading-order distributions are stationary, there will be no momentum transfer due to collisions at leading order.

Equation (3.3.1) can be solved for $\boldsymbol{u}_\perp^1$ to obtain

$$\boldsymbol{u}_\perp^1 = \frac{\boldsymbol{B} \times \nabla_\perp p_\perp^0}{n_\perp \omega_c \tau Z |B|^2} + \frac{\boldsymbol{E}_\perp \times \boldsymbol{B}}{\omega_c \tau |B|^2}, \tag{3.3.2}$$

which is recognizable as the sum of the diamagnetic and $E \times B$ drift velocities.

Having obtained $\boldsymbol{u}_\perp^1$, we substitute $\boldsymbol{u}_\perp = \epsilon \boldsymbol{u}_\perp^1 + \mathcal{O}(\epsilon^2)$ into (3.2.9) and (3.2.11) and retain terms up to order ϵ. This gives

$$\begin{aligned} &\frac{\mathrm{d}}{\mathrm{d}t^0} n_\perp + \epsilon \partial_{t^1} n_\perp + \epsilon \nabla_\perp \cdot (n_\perp \boldsymbol{u}_\perp^1) = \mathcal{N}^0 + \epsilon \mathcal{N}^1 \\ &\frac{\mathrm{d}}{\mathrm{d}t^0} e_\perp + \epsilon \partial_{t^1} e_\perp + \epsilon \nabla_\perp \cdot (2 e_\perp \mathbb{I} \cdot \boldsymbol{u}_\perp^1 + \boldsymbol{q}_\perp^1) = \epsilon n_\perp Z \boldsymbol{E}_\perp \cdot \boldsymbol{u}_\perp^1 + \mathcal{Q}^0 + \epsilon \mathcal{Q}^1. \end{aligned} \tag{3.3.3}$$

We have used the fact that $\mathbb{P}^0 = e_\perp \mathbb{I}$ and that $\boldsymbol{q}_\perp^0 = 0$.

The only flux term in (3.3.3) that requires closure is $\boldsymbol{q}_\perp^1$. Expanding the definition of heat flux and collecting order ϵ terms gives

$$\begin{aligned}
\boldsymbol{q}_\perp &= \frac{A}{2}\int(\boldsymbol{v}_\perp - \boldsymbol{u}_\perp)|\boldsymbol{v}_\perp - \boldsymbol{u}_\perp|^2 f \,\mathrm{d}\boldsymbol{v}_\perp \\
&= -\frac{\epsilon A}{2}\boldsymbol{u}_\perp^1\int|\boldsymbol{v}_\perp|^2 f^0\,\mathrm{d}\boldsymbol{v}_\perp - \frac{\epsilon A}{2}\int 2(\boldsymbol{v}_\perp\otimes\boldsymbol{v}_\perp)\cdot\boldsymbol{u}_\perp^1 f^0\,\mathrm{d}\boldsymbol{v}_\perp + \frac{\epsilon A}{2}\int\boldsymbol{v}_\perp|\boldsymbol{v}_\perp|^2 f^1\,\mathrm{d}\boldsymbol{v}_\perp + \mathcal{O}(\epsilon^2) \\
&= -\epsilon A\boldsymbol{u}_\perp^1\int|\boldsymbol{v}_\perp|^2 f^0\,\mathrm{d}\boldsymbol{v}_\perp + \frac{\epsilon A}{2}\int\boldsymbol{v}_\perp|\boldsymbol{v}_\perp|^2 f^1\,\mathrm{d}\boldsymbol{v}_\perp + \mathcal{O}(\epsilon^2) \\
&= \underbrace{-2\epsilon p_\perp^0\boldsymbol{u}_\perp^1 + \frac{\epsilon A}{2}\int\boldsymbol{v}_\perp|\boldsymbol{v}_\perp|^2 f^1\,\mathrm{d}\boldsymbol{v}_\perp}_{\boldsymbol{q}_\perp^1} + \mathcal{O}(\epsilon^2).
\end{aligned} \tag{3.3.4}$$

Evidently we must determine an expression for f^1, which can be done by solving the kinetic equation at order ϵ.

3.3.2 First-order correction

The equation for f^1 is

$$-\Omega_c\partial_\phi f^1 = -\partial_{t^0}\mathcal{M} - \boldsymbol{v}_\perp\cdot\nabla_\perp\mathcal{M} - \boldsymbol{E}_\perp\cdot\nabla_{\boldsymbol{v}_\perp}\mathcal{M} - \mathcal{D}_\parallel\mathcal{M} + C^0(f), \tag{3.3.5}$$

where we have used the azimuthal velocity coordinate ϕ which is defined via

$$\boldsymbol{v}_\perp = (v_\perp\cos\phi, v_\perp\sin\phi)^T,$$

and defined the cyclotron frequency for the current species as $\Omega_c = \omega_c\tau\frac{Z|B|}{A}$. The solution of (3.3.5) relies on a Fredholm compatibility condition on the right-hand side, which is that its gyroaverage should vanish. To express these constraints, we introduce notation to split a term g into its gyro-averaged component $\overline{g}$ and the remainder $\widetilde{g}$:

$$g = \overline{g} + \widetilde{g} = \frac{1}{2\pi}\int_0^{2\pi} g(\phi')\,\mathrm{d}\phi' + \left(g - \frac{1}{2\pi}\int_0^{2\pi} g(\phi')\,\mathrm{d}\phi'\right). \tag{3.3.6}$$

The Fredholm compatibility condition for (3.3.5) is then

$$\partial_{t^0}\mathcal{M} + \mathcal{D}_\parallel\mathcal{M} = \overline{C^0(f)}. \tag{3.3.7}$$

We return to our remark about assuming a Maxwellian zeroth-order solution. If instead we assume that $\partial_{t^0} = 0$, which is to say, assume that dynamics are on the whole slower than the collisional timescale, then we also obtain the result that the zeroth-order solution is a Maxwellian in the perpendicular direction. That is, assuming that $\partial_{t^0} = 0$ acts as a retroactive justification for our assumption that $f^0 = \mathcal{M}$.

Subtracting (3.3.7) from (3.3.5), the first-order equation is

$$\begin{aligned} -\Omega_c \partial_\phi f^1 &= -\boldsymbol{v}_\perp \cdot \nabla_\perp \mathcal{M} - \frac{Z}{A} \boldsymbol{E}_\perp \cdot \nabla_{\boldsymbol{v}_\perp} \mathcal{M} + \widetilde{C^0(f)} && (3.3.8) \\ &= -\boldsymbol{v}_\perp \cdot \left[\frac{\nabla n_\perp}{n_\perp} - \frac{\nabla T_\perp^0}{T_\perp^0} + \frac{A |\boldsymbol{v}_\perp|^2 \nabla T_\perp^0}{2 (T_\perp^0)^2} \right] \mathcal{M} + \frac{Z(\boldsymbol{E}_\perp \cdot \boldsymbol{v}_\perp)}{T_\perp^0}. && (3.3.9) \end{aligned}$$

We argue that the agyrotropic component of $C^0(f)$ vanishes just as we argued that $\mathcal{S}^0$ vanished. Since all of the zeroth-order distribution functions which participate in $C^0(f)$ are gyrotropic, this must be true by symmetry.

It is not hard to show that for a vector $\boldsymbol{g}$ which is independent of ϕ, that

$$\int \boldsymbol{v} \cdot \boldsymbol{g} \, \mathrm{d}\phi = \boldsymbol{g} \cdot \int \boldsymbol{v} \, \mathrm{d}\phi = \frac{(\boldsymbol{v} \times \boldsymbol{B}) \cdot \boldsymbol{g}}{|B|} = \frac{\boldsymbol{v} \cdot (\boldsymbol{B} \times \boldsymbol{g})}{|B|}$$

Thus,

$$\begin{aligned} f^1 &= \frac{-\boldsymbol{v}_\perp}{|B|\Omega_c} \cdot \boldsymbol{B} \times \left[-\frac{\nabla_\perp n_\perp}{n_\perp} + \frac{\nabla_\perp T_\perp^0}{T_\perp^0} + \frac{Z \boldsymbol{E}_\perp}{T_\perp^0} \right] \mathcal{M} + \frac{A |\boldsymbol{v}_\perp|^2 \boldsymbol{v}_\perp}{2|B|\Omega_c} \cdot \frac{\boldsymbol{B} \times \nabla_\perp T_\perp^0}{T_\perp^2} \mathcal{M} \\ &= \frac{\boldsymbol{v}_\perp}{|B|\Omega_c} \cdot \left[\frac{\boldsymbol{B} \times \nabla_\perp p_\perp^0}{p_\perp^0} + \frac{Z \boldsymbol{E} \times \boldsymbol{B}}{T_\perp^0} \right] \mathcal{M} - 2 \frac{\boldsymbol{v}_\perp \cdot \boldsymbol{B} \times \nabla_\perp T_\perp^0}{|B|\Omega_c T_\perp^0} \mathcal{M} + \frac{A |\boldsymbol{v}_\perp|^2 \boldsymbol{v}_\perp}{2|B|\Omega_c} \cdot \frac{\boldsymbol{B} \times \nabla_\perp T_\perp^0}{(T_\perp^0)^2} \mathcal{M} \\ &= \frac{A}{T_\perp^0} \left(\left[\left(\frac{\boldsymbol{B} \times \nabla_\perp p_\perp^0}{n_\perp \omega_c \tau Z |B|^2} + \frac{\boldsymbol{E} \times \boldsymbol{B}}{\omega_c \tau |B|^2} \right) - 2 \frac{\boldsymbol{B} \times \nabla_\perp T_\perp^0}{\omega_c \tau Z |B|^2} \right] \cdot \boldsymbol{v}_\perp \mathcal{M} + \frac{A}{T_\perp^0} \left[\frac{\boldsymbol{B} \times \nabla_\perp T_\perp^0}{2 \omega_c \tau Z |B|^2} \right] \cdot |\boldsymbol{v}_\perp|^2 \boldsymbol{v}_\perp \mathcal{M} \right) \\ &\triangleq \frac{A}{T_\perp^0} \left(\boldsymbol{u}_\perp^1 - 2 \boldsymbol{u}_T \right) \cdot \boldsymbol{v}_\perp \mathcal{M} + \frac{A^2}{(T_\perp^0)^2} \frac{\boldsymbol{u}_T}{2} \cdot \boldsymbol{v}_\perp |\boldsymbol{v}_\perp|^2 \mathcal{M}, && (3.3.10) \end{aligned}$$

where we have introduced the quantity $\boldsymbol{u}_T$ which has dimensions of velocity, given by

$$\boldsymbol{u}_T = \frac{\boldsymbol{B} \times \nabla_\perp T_\perp^0}{\omega_c \tau Z |B|^2}. \tag{3.3.11}$$

As we will see below, $\boldsymbol{u}_T$ is half the effective velocity at which thermal pressure is convected by the diamagnetic heat flux. It is also possible to view $\boldsymbol{u}_T$ as that part of the diamagnetic drift velocity attributable to temperature gradients, as opposed to density gradients.

Leading-order drift velocity

It is easy to verify that

$$n_\perp \boldsymbol{u}_\perp^1 = \int \boldsymbol{v}_\perp f^1 \, \mathrm{d}\boldsymbol{v}_\perp,$$

as we expect, by using properties of the Maxwellian. Indeed, using (3.3.10), we have

$$\begin{aligned}\int \boldsymbol{v}_\perp f^1 \,\mathrm{d}\boldsymbol{v}_\perp &= \frac{A}{T_\perp^0}(\boldsymbol{u}_\perp^1 - 2\boldsymbol{u}_T)\cdot\int(\boldsymbol{v}_\perp\otimes\boldsymbol{v}_\perp)\mathcal{M}\,\mathrm{d}\boldsymbol{v}_\perp + \frac{A^2}{(T_\perp^0)^2}\frac{\boldsymbol{u}_T}{2}\cdot\int(\boldsymbol{v}_\perp\otimes\boldsymbol{v}_\perp)|\boldsymbol{v}_\perp|^2\mathcal{M}\,\mathrm{d}\boldsymbol{v}_\perp \\ &= \frac{A}{T_\perp^0}(\boldsymbol{u}_\perp^1 - 2\boldsymbol{u}_T)\cdot\frac{p_\perp^0}{A}\mathbb{I} + \frac{A^2}{(T_\perp^0)^2}\frac{\boldsymbol{u}_T}{2}\cdot\frac{4p_\perp^0 T_\perp^0}{A^2}\mathbb{I} \\ &= n_\perp \boldsymbol{u}_\perp^1.\end{aligned} \tag{3.3.12}$$

Leading-order heat flux

Recalling that the first-order heat flux was defined in (3.3.4), the only integral remaining to calculate is the contracted third moment of f^1, which is as follows:

$$\begin{aligned}\frac{A}{2}\int \boldsymbol{v}_\perp|\boldsymbol{v}_\perp|^2 f^1\,\mathrm{d}\boldsymbol{v}_\perp &= \frac{A^2}{2T_\perp^0}(\boldsymbol{u}_\perp^1 - 2\boldsymbol{u}_T)\cdot\int(\boldsymbol{v}_\perp\otimes\boldsymbol{v}_\perp)|\boldsymbol{v}_\perp|^2\mathcal{M}\,\mathrm{d}\boldsymbol{v}_\perp + \frac{A^3}{2(T_\perp^0)^2}\frac{\boldsymbol{u}_T}{2}\cdot\int(\boldsymbol{v}_\perp\otimes\boldsymbol{v}_\perp)|\boldsymbol{v}_\perp|^4\mathcal{M}\,\mathrm{d}\boldsymbol{v}_\perp \\ &= \frac{4p_\perp^0\boldsymbol{u}_\perp^1}{2} - \frac{8p_\perp^0\boldsymbol{u}_T}{2} + \frac{24p_\perp^0\boldsymbol{u}_T}{4} \\ &= 2p_\perp^0\boldsymbol{u}_\perp^1 + 2p_\perp^0\boldsymbol{u}_T.\end{aligned} \tag{3.3.13}$$

Combining (3.3.13) and (3.3.4), we get an expression for $\boldsymbol{q}_\perp^1$:

$$\boldsymbol{q}_\perp^1 = 2p_\perp^0\boldsymbol{u}_T. \tag{3.3.14}$$

Leading-order pressure tensor

From the expression (3.3.10) for f^1, we can see immediately that it has odd-order polynomial dependence on $\boldsymbol{v}_\perp$. It follows that the even-order moments of f^1 vanish, and in particular the pressure tensor vanishes:

$$\mathbb{P}^1 = \int(\boldsymbol{v}_\perp - \boldsymbol{u}_\perp)\otimes(\boldsymbol{v}_\perp - \boldsymbol{u}_\perp)f^1\,\mathrm{d}\boldsymbol{v}_\perp = 0.$$

3.3.3 Drift-advection formulation

Having calculated $\mathbb{P}^1$ and $\boldsymbol{q}_\perp^1$, we can simplify our first-order semi-fluid equations considerably. First of all, note that $\nabla_\perp\cdot(n_\perp\boldsymbol{u}^1) = \nabla_\perp\cdot(n_\perp\boldsymbol{u}_E) = \boldsymbol{u}_E\cdot\nabla_\perp n_\perp$, where

$$\boldsymbol{u}_E = \frac{\boldsymbol{E}\times\boldsymbol{B}}{\omega_c\tau|B|^2}.$$

The continuity equation thereby decouples from the energy equation, so we have

$$\frac{\mathrm{d}}{\mathrm{d}t}n_\perp + \nabla_\perp\cdot(n_\perp\boldsymbol{u}_E) = \mathcal{N}^0 \tag{3.3.15}$$

Moreover, in the energy equation, we can calculate

$$\nabla_\perp \cdot (2p_\perp^0 \boldsymbol{u}_\perp^1 + 2p_\perp^0 \boldsymbol{u}_T) = \nabla_\perp \cdot \left(T_\perp^0 \frac{\boldsymbol{B} \times \nabla_\perp p_\perp^0}{\omega_c \tau Z |B|^2} + p_\perp^0 \frac{\boldsymbol{B} \times \nabla_\perp T_\perp^0}{\omega_c \tau Z |B|^2} \right) = \nabla_\perp \cdot \left(\frac{\boldsymbol{B} \times \nabla_\perp (p_\perp^0 T_\perp^0)}{\omega_c \tau Z |B|^2} \right) = 0.$$

Therefore, the energy equation simplifies to

$$\frac{\mathrm{d}}{\mathrm{d}t} e_\perp = n_\perp Z \boldsymbol{E} \cdot \boldsymbol{u}_\perp^1 + \mathcal{Q}^0.$$

Because of the decoupling between energy and continuity equations, it is possible to solve (3.3.15) only, assuming that collisional exchange of energy between particle species is neglected.

3.4 Hyperbolic expansion

In the hyperbolic expansion, we choose to retain $\boldsymbol{u}_\perp$ as a primary unknown, even though it does not appear as a spatially-dependent parameter in the zeroth-order solution $\mathcal{M}$. As we will see, rather than directly inverting the first-order equation, we may propose a solution f^1 which depends on $\boldsymbol{u}_\perp$, and satisfies the first-order correction equation to order ϵ^2. This allows us to retain a momentum equation.

3.4.1 Hyperbolic semi-fluid equations

In the preceding section we showed that the leading-order drift velocity is order ϵ:

$$\boldsymbol{u}_\perp = \epsilon \boldsymbol{u}_\perp^1 + \mathcal{O}(\epsilon^2).$$

We now seek a system of equations that treats $\boldsymbol{u}_\perp$ as a primary unknown, so that we may retain a momentum equation that captures some inertial effects. To that end, we rescale the drift velocity to

$$\boldsymbol{u}_\perp^* = \epsilon^{-1} \boldsymbol{u}_\perp.$$

The hyperbolic semi-fluid system of equations is obtained by retaining terms of up to order ϵ^2 in the semi-fluid equations, which we rewrite in terms of $\boldsymbol{u}_\perp^*$. The resulting system is

$$\frac{\mathrm{d}}{\mathrm{d}t} n_\perp + \epsilon \nabla_\perp (n_\perp \boldsymbol{u}_\perp^*) = \mathcal{N}^{0,1,2} \tag{3.4.1}$$

$$\epsilon A \frac{\mathrm{d}}{\mathrm{d}t} (n_\perp \boldsymbol{u}_\perp^*) + \nabla_\perp \cdot \left(\epsilon^2 A n_\perp \boldsymbol{u}_\perp^* \otimes \boldsymbol{u}_\perp^* + \mathbb{P}^{0,1,2} \right) = n_\perp Z (\boldsymbol{E} + \boldsymbol{u}_\perp^* \times \boldsymbol{B}) + \mathcal{S}^{0,1,2} \tag{3.4.2}$$

$$\frac{\mathrm{d}}{\mathrm{d}t} e_\perp + \nabla_\perp \cdot (\epsilon (e_\perp \mathbb{I} + \mathbb{P}^{0,1}) \cdot \boldsymbol{u}_\perp^* + \boldsymbol{q}_\perp^{0,1,2}) = \epsilon n_\perp Z \boldsymbol{E}_\perp \cdot \boldsymbol{u}_\perp^* + \mathcal{Q}^{0,1,2}. \tag{3.4.3}$$

The notation $\psi^{0,1,2}$ indicates the truncated Hilbert expansion

$$\psi^{0,1,2} = \psi^0 + \epsilon\psi^1 + \epsilon^2\psi^2.$$

The unclosed moments appearing in (3.4.2) and (3.4.3) are the familiar full pressure tensor $\mathbb{P}_\perp^{0,1,2}$ and the heat flux $\boldsymbol{q}_\perp^{0,1,2}$. We will obtain expressions for these via a Hilbert expansion procedure of the kinetic equation, much as we did with the diffusive expansion.

3.4.2 Modified first-order correction

To obtain closures for the hyperbolic semi-fluid equations in terms of $\boldsymbol{u}_\perp$, we must manipulate the kinetic equation slightly. The hierarchy of kinetic equations, collecting terms of equal order in ϵ, is

$$-\epsilon^{-1}\Omega_c\partial_\phi\mathcal{M} = 0 \tag{3.4.4}$$

$$-\Omega_c\partial_\phi f^1 = -\frac{\mathrm{d}}{\mathrm{d}t^0}\mathcal{M} - \boldsymbol{v}_\perp\cdot\nabla_\perp\mathcal{M} - \boldsymbol{E}_\perp\cdot\nabla_{\boldsymbol{v}_\perp}\mathcal{M} + C^0(f) \tag{3.4.5}$$

$$-\epsilon\Omega_c\partial_\phi f^2 = \epsilon\left(-\partial_{t^1}\mathcal{M} - \frac{\mathrm{d}}{\mathrm{d}t^0}f^1 - \boldsymbol{v}_\perp\cdot\nabla_\perp f^1 - \boldsymbol{E}_\perp\cdot\nabla_{\boldsymbol{v}_\perp}f^1 + C^1(f)\right) \tag{3.4.6}$$

$$\vdots \tag{3.4.7}$$

Here, each subsequent kinetic solution $\mathcal{M}, f^1, f^2, \ldots$ is the solution as calculated by the diffusive expansion. We now collect all of the drift velocity into the velocity of the first-order solution. All second-order and higher kinetic solutions will have zero drift velocity. To do so, we add the following (vanishing) sum, term by term, to the right-hand side of (3.4.4):

$$-\Omega_c\partial_\phi\frac{A}{T_\perp^0}\left(\boldsymbol{u}_\perp^* - \boldsymbol{u}_\perp^1 - \epsilon\boldsymbol{u}_\perp^2 - \epsilon^2\boldsymbol{u}_\perp^3 - \cdots\right)\cdot\boldsymbol{v}_\perp\mathcal{M} = 0,$$

where $\boldsymbol{u}_\perp^1, \boldsymbol{u}_\perp^2, \ldots$ are to be computed using successive orders of the semi-fluid momentum equation, as in (3.3.2).

To distinguish kinetic solutions of the modified equation hierarchy, we introduce the following "hatted" notation, in terms of which the kinetic equation is, at each order of ϵ:

$$\begin{aligned}
-\epsilon^{-1}\Omega_c\partial_\phi\mathcal{M} &= 0\\
-\Omega_c\partial_\phi\hat{f}^1 &= -\frac{\mathrm{d}}{\mathrm{d}t^0}\mathcal{M} - \boldsymbol{v}_\perp\cdot\nabla_\perp\mathcal{M} - \boldsymbol{E}_\perp\cdot\nabla_{\boldsymbol{v}_\perp}\mathcal{M} + C^0(f) - \Omega_c\partial_\phi\frac{A}{T_\perp^0}(\boldsymbol{u}_\perp^* - \boldsymbol{u}_\perp^1)\cdot\boldsymbol{v}_\perp\mathcal{M}\\
-\epsilon\Omega_c\partial_\phi\hat{f}^2 &= \epsilon\left(-\partial_{t^1}\mathcal{M} - \frac{\mathrm{d}}{\mathrm{d}t^0}\hat{f}^1 - \boldsymbol{v}_\perp\cdot\nabla_\perp\hat{f}^1 - \boldsymbol{E}_\perp\cdot\nabla_{\boldsymbol{v}_\perp}\hat{f}^1 + C^1(f)\right) + \epsilon\Omega_c\partial_\phi\frac{A}{T_\perp^0}\boldsymbol{u}_\perp^2\cdot\boldsymbol{v}_\perp\mathcal{M}\\
&\vdots
\end{aligned}$$

Following a similar argument as in Section 3.3, the first-order kinetic equation has the solution

$$\hat{f}^1 = \frac{A}{T_\perp^0}(\boldsymbol{u}_\perp^* - 2\boldsymbol{u}_T)\cdot \boldsymbol{v}_\perp \mathcal{M} + \frac{A^2}{(T_\perp^0)^2}\frac{\boldsymbol{u}_T}{2}\cdot \boldsymbol{v}_\perp |\boldsymbol{v}_\perp|^2 \mathcal{M}. \tag{3.4.8}$$

The fluid velocity of the first-order solution is

$$\int \boldsymbol{v}_\perp \hat{f}^1 \,\mathrm{d}\boldsymbol{v}_\perp = n_\perp \boldsymbol{u}_\perp^*. \tag{3.4.9}$$

By adding terms that sum to 0 to (3.4.4), we did not change the underlying kinetic equation which the Hilbert expansion of f satisfies. We merely rearranged terms to locate all momentum in $\hat{f}^1$. Therefore, the hatted solutions form an equally valid expansion of f:

$$f = \mathcal{M} + \epsilon \hat{f}^1 + \epsilon^2 \hat{f}^2 + \cdots \tag{3.4.10}$$

Leading-order heat flux

A similar calculation to (3.3.13) shows that

$$\frac{A}{2}\int \boldsymbol{v}_\perp |\boldsymbol{v}_\perp|^2 \hat{f}^1 \,\mathrm{d}\boldsymbol{v}_\perp = 2p_\perp^0 \boldsymbol{u}_\perp^* + 2p_\perp^0 \boldsymbol{u}_T,$$

which shows that the leading-order heat flux closure is unchanged from the diffusive expansion:

$$\boldsymbol{q}_\perp^1 = 2p_\perp^0 \boldsymbol{u}_T.$$

Full pressure tensor

We now seek an expression for the full pressure tensor. Expanding the definition (3.2.3), we have

$$\begin{aligned}
\mathbb{P}_\perp &= A\int (\boldsymbol{v}_\perp \otimes \boldsymbol{v}_\perp) f \,\mathrm{d}\boldsymbol{v}_\perp \\
&= A\int (\boldsymbol{v}_\perp \otimes \boldsymbol{v}_\perp)\mathcal{M}\,\mathrm{d}\boldsymbol{v}_\perp + A\int (\boldsymbol{u}_\perp \otimes \boldsymbol{u}_\perp)\mathcal{M}\,\mathrm{d}\boldsymbol{v}_\perp \\
&\quad - A\epsilon \int (\boldsymbol{v}_\perp \otimes \boldsymbol{u}_\perp)\hat{f}^1 \,\mathrm{d}\boldsymbol{v}_\perp - A\epsilon \int (\boldsymbol{u}_\perp \otimes \boldsymbol{v}_\perp)\hat{f}^1 \,\mathrm{d}\boldsymbol{v}_\perp \\
&\quad + A\epsilon^2 \int (\boldsymbol{v}_\perp \otimes \boldsymbol{v}_\perp)\hat{f}^2 \,\mathrm{d}\boldsymbol{v}_\perp + \mathcal{O}(\epsilon^3) \\
&= p_\perp^0 \mathbb{I} - \epsilon^2 A n_\perp (\boldsymbol{u}_\perp^* \otimes \boldsymbol{u}_\perp^*) + \epsilon^2 A \int (\boldsymbol{v}_\perp \otimes \boldsymbol{v}_\perp)\hat{f}^2 \,\mathrm{d}\boldsymbol{v}_\perp + \mathcal{O}(\epsilon^3),
\end{aligned} \tag{3.4.11}$$

where we have used (3.4.9). Evidently, the trace-free components of $\mathbb{P}_\perp$ are of order ϵ^2 at least, and we must calculate the full second moment tensor of $\hat{f}^2$ to obtain a complete expression for $\mathbb{P}^{0,1,2}$.

3.4.3 Second-order correction

The modified second-order equation is

$$-\Omega_c \partial_\phi \hat{f}^2 = -\partial_{t^1}\mathcal{M} - \partial_{t^0}\hat{f}^1 - \boldsymbol{v}_\perp \cdot \nabla_\perp \hat{f}^1 - \frac{Z}{A}\boldsymbol{E}_\perp \cdot \nabla_{\boldsymbol{v}_\perp}\hat{f}^1 - \mathcal{D}_\parallel \hat{f}^1 + C^1(f) + \Omega_c \partial_\phi \frac{A}{T_\perp^0}\boldsymbol{u}_\perp^2 \cdot \boldsymbol{v}_\perp \mathcal{M} \quad (3.4.12)$$

Again there is a Fredholm compatibility condition on the right-hand side of this equation, that its gyroaverage vanish. It is up to the $\partial_{t^1}\mathcal{M}$ term to eliminate all of the gyrotropic components of the other terms on the right-hand side. To do so, we take the zeroth and second moments of (3.4.12) to obtain semi-fluid equations satisfied by $n_\perp$ and $T_\perp^0$:

$$\partial_{t^1} n_\perp + \nabla_\perp \cdot (n_\perp \boldsymbol{u}_\perp^*) = \mathcal{N}^1 \quad (3.4.13)$$

$$\partial_{t^1} p_\perp^0 + \nabla_\perp \cdot (2p_\perp^0 \boldsymbol{u}_\perp^* + \boldsymbol{q}_\perp^1) = n_\perp Z \boldsymbol{E}_\perp \cdot \boldsymbol{u}_\perp^* + \mathcal{Q}^1. \quad (3.4.14)$$

Equations (3.4.13) and (3.4.14) can then be used to eliminate the $\partial_{t^1}\mathcal{M}$ term. After expanding all gyrotropic terms in (3.4.12) and simplifying, we find that the gyrotropic terms amount to

$$\overline{\partial_{t^1}\mathcal{M} + \boldsymbol{v}_\perp \cdot \nabla_\perp \hat{f}^1 + \frac{Z}{A}\boldsymbol{E}_\perp \cdot \nabla_{\boldsymbol{v}_\perp}\hat{f}^1} = \left(2\frac{\mathcal{N}^1}{n_\perp} - \frac{\mathcal{Q}^1}{p_\perp^0} + \frac{A|\boldsymbol{v}_\perp|^2}{2T}\left(\frac{\mathcal{Q}^1}{p_\perp^0} - \frac{\mathcal{N}^1}{n_\perp}\right)\right)\mathcal{M} + \mathcal{O}(\epsilon).$$

The Fredholm condition on (3.4.12) will therefore be satisfied if $\mathcal{N}^1 = \mathcal{Q}^1 = 0$. This is quite a natural condition, since $\mathcal{N}$ and $\mathcal{Q}$ are gyrotropic moments of the collision operator, whereas $\hat{f}^1$ is agyrotropic.

The remaining (agyrotropic) terms of Equation (3.4.12) are as follows:

$$-\Omega_c \partial_\phi \hat{f}^2 = -\boldsymbol{v}_\perp \widetilde{\cdot \nabla_\perp}\hat{f}^1 - \frac{Z}{A}\boldsymbol{E}_\perp \widetilde{\cdot \nabla_{\boldsymbol{v}_\perp}}\hat{f}^1 - \frac{\mathrm{d}}{\mathrm{d}t^0}\hat{f}^1 - \widetilde{C^1(f)} + \Omega_c \partial_\phi \frac{A}{T_\perp^0}\boldsymbol{u}_\perp^2 \cdot \boldsymbol{v}_\perp \mathcal{M}, \quad (3.4.15)$$

where the tilde notation indicates the agyrotropic part of the term, as defined in (3.3.6). Note that $\frac{\mathrm{d}}{\mathrm{d}t^0}\hat{f}^1$ is a fully agyrotropic term already. We are not interested in an explicit expression for $\hat{f}^2$, but rather in its moments. Thus, we will not invert this whole equation, but rather consider each term and the contribution that it makes to the moments of $\hat{f}^2$. The first two terms on the right-hand side of (3.4.15) have a polynomial dependence on $\boldsymbol{v}_\perp$ which is of even total degree, so they will make contributions to the pressure tensor. The latter trio of terms are of odd total order in $\boldsymbol{v}_\perp$ and therefore make contributions to the heat flux, but not the pressure tensor.

Full pressure tensor

As shown in (3.4.11), to complete the expression for $\mathbb{P}^{0,1,2}$ we must calculate the following moment of $\hat{f}^2$:

$$A \int (\boldsymbol{v}_\perp \otimes \boldsymbol{v}_\perp) \hat{f}^2$$

It will be convenient to write $\hat{f}^1$ in a more abstract form which highlights its velocity dependence:

$$f^1 = \boldsymbol{a} \cdot \boldsymbol{v}_\perp \mathcal{M} + \boldsymbol{b} \cdot \boldsymbol{v}_\perp |\boldsymbol{v}_\perp|^2 \mathcal{M}, \tag{3.4.16}$$

where

$$\boldsymbol{a} = \frac{A}{T_\perp^0}(\boldsymbol{u}_\perp^* - 2\boldsymbol{u}_T), \quad \boldsymbol{b} = \frac{A^2}{(T_\perp^0)^2} \frac{\boldsymbol{u}_T}{2}. \tag{3.4.17}$$

We also introduce the pair of tensors

$$\begin{aligned}
\mathbb{U} = (\widetilde{\boldsymbol{v}_\perp \otimes \boldsymbol{v}_\perp}) &= \begin{pmatrix} v_x^2 - \frac{|\boldsymbol{v}_\perp|^2}{2} & v_x v_y \\ v_y v_x & v_y^2 - \frac{|\boldsymbol{v}_\perp|^2}{2} \end{pmatrix}, \\
\mathbb{U}^* = \int \mathbb{U} \, \mathrm{d}\phi &= \begin{pmatrix} \frac{v_x v_y}{2} - \frac{1}{4} & \frac{v_y^2 - v_x^2}{4} & 0 \\ \frac{v_y^2 - v_x^2}{4} & \frac{1}{4} - \frac{v_x v_y}{2} & 0 \\ 0 & 0 & 0 \end{pmatrix},
\end{aligned}$$

where we have denoted the two components of $\boldsymbol{v}_\perp$ by v_x, v_y for purposes of illustration. The terms with contributions to the pressure tensor are then given by

$$\begin{aligned}
\frac{1}{\omega_c} \int \left[\widetilde{\boldsymbol{v}_\perp \cdot \nabla_\perp} f^1 + \frac{Z}{A} \widetilde{\boldsymbol{E}_\perp \cdot \nabla_{\boldsymbol{v}_\perp}} f^1 \right] \mathrm{d}\phi &= [(\nabla_\perp \boldsymbol{a})\mathcal{M} + \boldsymbol{a}(\nabla_\perp \mathcal{M})] : \mathbb{U}^* + [(\nabla_\perp \boldsymbol{b})\mathcal{M} + \boldsymbol{b}(\nabla_\perp \mathcal{M})] : \mathbb{U}^* |\boldsymbol{v}_\perp|^2 \\
&\quad + \frac{Z}{A}(\boldsymbol{E} \otimes \boldsymbol{a}) : \left(-\frac{A}{T_\perp^0} \mathbb{U}^* \right) \mathcal{M} + \frac{Z}{A}(\boldsymbol{E} \otimes \boldsymbol{b}) : \left(2\mathbb{U}^* - \frac{A|\boldsymbol{v}_\perp|^2}{T_\perp^0} \mathbb{U}^* \right) \mathcal{M} \\
&\triangleq G.
\end{aligned}$$

Recalling from (3.2.18) that

$$\begin{aligned}
\nabla_\perp \mathcal{M} &= \left(\frac{\nabla_\perp n_\perp}{n_\perp} - \frac{\nabla_\perp T_\perp^0}{T_\perp^0} + \frac{A|\boldsymbol{v}_\perp|^2 \nabla_\perp T_\perp^0}{2(T_\perp^0)^2} \right) \mathcal{M} \\
&\triangleq \boldsymbol{p}\mathcal{M} + \frac{A|\boldsymbol{v}_\perp|^2}{T_\perp^0} \boldsymbol{r}\mathcal{M},
\end{aligned}$$

we can expand G further:

$$\begin{aligned}G = &\left[(\nabla_\perp \boldsymbol{a}) + \boldsymbol{a}\otimes\boldsymbol{p} + \frac{A|\boldsymbol{v}_\perp|^2}{T_\perp^0}\boldsymbol{a}\otimes\boldsymbol{r} + (\nabla_\perp \boldsymbol{b})|\boldsymbol{v}_\perp|^2 + \boldsymbol{b}\otimes\boldsymbol{p}|\boldsymbol{v}_\perp|^2 + \frac{A|\boldsymbol{v}_\perp|^4}{T_\perp^0}\boldsymbol{b}\otimes\boldsymbol{r}\right]\mathcal{M} : \mathbb{U}^* \\ &+ \frac{Z}{A}\left[-\frac{A}{T_\perp^0}\boldsymbol{E}\otimes\boldsymbol{a} + 2\boldsymbol{E}\otimes\boldsymbol{b} - \frac{A|\boldsymbol{v}_\perp|^2}{T_\perp^0}\boldsymbol{E}\otimes\boldsymbol{b}\right]\mathcal{M} : \mathbb{U}^*\end{aligned}$$

Computing the pressure of this expression can be done with a pair of useful identities for the quadratic moments of $\mathbb{U}^*\mathcal{M}$, which we summarize here:

$$\int |\boldsymbol{v}_\perp|^i \frac{v_x^2 - v_y^2}{2}\mathbb{U}^*\mathcal{M}\,\mathrm{d}\boldsymbol{v} = \frac{p_\perp^0 (T_\perp^0)^{i+1}}{A^{i+2}}\begin{pmatrix} 0 & -c(i) & 0 \\ -c(i) & 0 & 0 \\ 0 & 0 & 0\end{pmatrix}$$

$$\int |\boldsymbol{v}_\perp|^i v_x v_y \mathbb{U}^*\mathcal{M}\,\mathrm{d}\boldsymbol{v} = \frac{p_\perp^0 (T_\perp^0)^{i+1}}{A^{i+2}}\begin{pmatrix} c(i) & 0 & 0 \\ 0 & -c(i) & 0 \\ 0 & 0 & 0\end{pmatrix}$$

The function $c(i)$ takes the values $c(0) = 1/2$, $c(2) = 3$, and $c(4) = 24$. Using these properties, we can derive the following useful identity for general vectors $\boldsymbol{\xi}$, $\boldsymbol{\zeta}$, and their outer product:

$$\begin{aligned}(\boldsymbol{\xi}\otimes\boldsymbol{\zeta}) : \int |\boldsymbol{v}_\perp|^i (\widetilde{\boldsymbol{v}_\perp \otimes \boldsymbol{v}_\perp})\mathbb{U}^*\mathcal{M}\,\mathrm{d}\boldsymbol{v} &= c(i)\frac{p_\perp^0 (T_\perp^0)^{i+1}}{A^{i+2}}\begin{pmatrix} -\xi_x\zeta_y - \xi_y\zeta_x & \xi_x\zeta_x - \xi_y\zeta_y \\ \xi_x\zeta_x - \xi_y\zeta_y & \xi_x\zeta_y + \xi_y\zeta_x\end{pmatrix} \\ &\triangleq c(i)\frac{p_\perp^0 (T_\perp^0)^{i+1}}{A^{i+2}}(\boldsymbol{\xi}\star\boldsymbol{\zeta})\end{aligned} \tag{3.4.18}$$

Here the double contraction is taken between $(\boldsymbol{\xi}\otimes\boldsymbol{\zeta})$ and the indices of $\mathbb{U}^*$. Note that the definition (3.4.18) also encompasses the case where $\boldsymbol{\xi} = \nabla_\perp$.

Now we are in a position to express the pressure tensor of $\hat{f}^2$ in terms of the $\star$ operator defined by (3.4.18):

$$
\begin{aligned}
A\int(\boldsymbol{v}_\perp\otimes\boldsymbol{v}_\perp)\hat{f}^2\,\mathrm{d}\boldsymbol{v}_\perp &= A\int(\boldsymbol{v}_\perp\otimes\boldsymbol{v}_\perp)G\,\mathrm{d}\boldsymbol{v}_\perp \\
&= \frac{A}{2}\frac{p^0_\perp T^0_\perp}{A^2}\left(\nabla_\perp\star\boldsymbol{a}+\boldsymbol{a}\star\boldsymbol{p}-\frac{Z}{T^0_\perp}\boldsymbol{E}\star\boldsymbol{a}+\frac{2Z}{A}\boldsymbol{E}\star\boldsymbol{b}\right) \\
&\quad+3A\frac{p^0_\perp(T^0_\perp)^2}{A^3}\left(\frac{A}{T^0_\perp}\boldsymbol{a}\star\boldsymbol{r}+\nabla_\perp\star\boldsymbol{b}+\boldsymbol{b}\star\boldsymbol{p}-\frac{Z}{T^0_\perp}\boldsymbol{E}\star\boldsymbol{b}\right) \\
&\quad+24A\frac{p^0_\perp(T^0_\perp)^3}{A^4}\frac{A}{T^0_\perp}\boldsymbol{b}\star\boldsymbol{r} \\
&= \frac{p^0_\perp T^0_\perp}{A}\nabla_\perp\star\left(\frac{\boldsymbol{a}}{2}+\frac{3T^0_\perp}{A}\boldsymbol{b}\right)+\frac{p^0_\perp T^0_\perp}{A}\left(\frac{\boldsymbol{a}}{2}+\frac{3T^0_\perp}{A}\boldsymbol{b}\right)\star\boldsymbol{p} \\
&\quad+\frac{Zp^0_\perp}{A}\boldsymbol{E}\star\left(\frac{-\boldsymbol{a}}{2}-\frac{2T^0_\perp}{A}\boldsymbol{b}\right)+\frac{p^0_\perp T^0_\perp}{A}\left(3\boldsymbol{a}+\frac{24T^0_\perp}{A}\boldsymbol{b}\right)\star\boldsymbol{r}
\end{aligned}
$$

Substituting $\boldsymbol{a}$ and $\boldsymbol{b}$ from their definitions (3.4.17),

$$
A\int(\boldsymbol{v}_\perp\otimes\boldsymbol{v}_\perp)\hat{f}^2\,\mathrm{d}\boldsymbol{v}_\perp = p^0_\perp\left[\nabla_\perp\star\frac{\boldsymbol{u}^*_\perp+\boldsymbol{u}_T}{2}+\frac{\boldsymbol{u}^*_\perp+\boldsymbol{u}_T}{2}\star\boldsymbol{p}\right]+\frac{Zp^0_\perp}{T^0_\perp}\boldsymbol{E}\star\frac{-\boldsymbol{u}^*_\perp}{2}+p^0_\perp(3\boldsymbol{u}^*_\perp+6\boldsymbol{u}_T)\star\boldsymbol{r}
$$

After some effort, it can be shown that the second moment of $\hat{f}^2$ is

$$
A\int\boldsymbol{v}_\perp\otimes\boldsymbol{v}_\perp\hat{f}^2\,\mathrm{d}\boldsymbol{v}_\perp = p^0_\perp\frac{\mathbb{W}_3[\boldsymbol{u}^*_\perp]}{2\Omega_c}+\frac{\mathbb{W}_3[\boldsymbol{q}^1_\perp]}{4\Omega_c}+\frac{A}{T^0_\perp}\left(\frac{p^0_\perp(\widehat{\boldsymbol{u}^1_\perp\otimes\boldsymbol{u}^*_\perp})}{2}+\frac{(\widehat{\boldsymbol{q}^1_\perp\otimes\boldsymbol{u}^*_\perp})}{4}+\frac{(\widehat{\boldsymbol{q}^1_\perp\otimes\boldsymbol{q}^1_\perp})}{2p^0_\perp}\right),
$$

where the symmetric, trace-free tensor $\mathbb{W}_3[\boldsymbol{u}]$ for a vector $\boldsymbol{u}=(u_x,u_y)$ is given by

$$
\epsilon_{xy\|}\begin{pmatrix}-\partial_y u_x-\partial_x u_y & \partial_x u_x-\partial_y u_y \\ \partial_x u_x-\partial_y u_y & \partial_y u_x+\partial_x u_y\end{pmatrix},
$$

with $\epsilon_{xy\|}$ a Levi-Civita symbol indicating the orientation of the triplet $(x,y,\|)$ for perpendicular coordinates x,y.

The notation $\widehat{\mathbb{A}}$ is defined as the trace-free symmetrization of $\mathbb{A}$:

$$
\widehat{\mathbb{A}}=\mathbb{A}+\mathbb{A}^T-\mathrm{Tr}(\mathbb{A})\mathbb{I}.
$$

These calculations lead us to our final expression for the pressure tensor up to order ϵ^2, which is

$$
\mathbb{P}^{0,1,2}=p^0_\perp\mathbb{I}-\epsilon^2An_\perp(\boldsymbol{u}^*_\perp\otimes\boldsymbol{u}^*_\perp)+\epsilon^2\left[p^0_\perp\frac{\mathbb{W}_3[\boldsymbol{u}^*_\perp]}{2\Omega_c}+\frac{\mathbb{W}_3[\boldsymbol{q}^1_\perp]}{4\Omega_c}+\frac{A}{T^0_\perp}\left(\frac{p^0_\perp(\widehat{(\boldsymbol{u}^1_\perp+\boldsymbol{u}_T)\otimes\boldsymbol{u}^*_\perp})}{2}+\frac{(\widehat{\boldsymbol{q}^1_\perp\otimes\boldsymbol{q}^1_\perp})}{2p^0_\perp}\right)\right]. \tag{3.4.19}
$$

Stress tensor

It is more useful to recombine our expression for the full pressure tensor $\mathbb{P}^{0,1,2}$ in terms of a scalar pressure and stress tensor. We regroup the first two terms and define $\boldsymbol{u}_\perp^\dagger = \boldsymbol{u}_\perp^1 + \boldsymbol{u}_T$ and $p_\perp = p_\perp^0 - An_\perp \frac{|\boldsymbol{u}_\perp|^2}{2}$. This gives

$$\mathbb{P}^{0,1,2} = p_\perp \mathbb{I} + \epsilon^2 \Pi_\perp^2,$$

where

$$\Pi_\perp^2 = p_\perp^0 \frac{\mathbb{W}_3[\boldsymbol{u}_\perp^*]}{2\Omega_c} + \frac{\mathbb{W}_3[\boldsymbol{q}_\perp^1]}{4\Omega_c} + \frac{An_\perp}{2}[\widehat{(\boldsymbol{u}_\perp^\dagger - \boldsymbol{u}_\perp^*) \otimes \boldsymbol{u}_\perp^*}] + \frac{A}{2p_\perp^0 T_\perp^0}(\widehat{\boldsymbol{q}_\perp^1 \otimes \boldsymbol{q}_\perp^1}).$$

Second-order heat flux in the absence of collisions and parallel dynamics

The second-order correction to the heat flux can be calculated by collecting order ϵ^2 terms in the definition of $\boldsymbol{q}_\perp$, which gives

$$\begin{aligned} \boldsymbol{q}_\perp^2 = \frac{A}{2} \int \boldsymbol{v}_\perp |\boldsymbol{u}_\perp^1|^2 f^0 \, \mathrm{d}\boldsymbol{v}_\perp - A \int \boldsymbol{v}_\perp \cdot (\boldsymbol{u}_\perp^1 \otimes \boldsymbol{u}_\perp^1) f^0 \, \mathrm{d}\boldsymbol{v}_\perp + \frac{A}{2} \int |\boldsymbol{v}_\perp|^2 \boldsymbol{u}_\perp^2 f^0 \, \mathrm{d}\boldsymbol{v}_\perp \\ + A \int \boldsymbol{u}_\perp^1 \cdot (\boldsymbol{v}_\perp \otimes \boldsymbol{v}_\perp) f^1 \, \mathrm{d}\boldsymbol{v}_\perp + \frac{A}{2} \int \boldsymbol{v}_\perp |\boldsymbol{v}_\perp|^2 f^2 \, \mathrm{d}\boldsymbol{v}_\perp. \end{aligned} \tag{3.4.20}$$

Odd-order moments of f^0 and even-order moments of f^1 vanish because f^0 and f^1 are respectively gyrotropic and agyrotropic in $\boldsymbol{v}_\perp$. Retaining terms of order ϵ in the semi-fluid momentum equation, we obtain an equation for $\boldsymbol{u}_\perp^2$:

$$A \frac{\mathrm{d}}{\mathrm{d}t^0}(n_\perp \boldsymbol{u}_\perp^1) = n_\perp Z \omega_c \tau \boldsymbol{u}_\perp^2 \times \boldsymbol{B} + \mathcal{S}^1,$$

where the pressure tensor term vanishes because $\mathbb{P}^1 = 0$. Assuming that $C(f) = 0$ the collisional term $\mathcal{S}^1$ vanishes. If we further assume that the plasma is uniform in the parallel direction, we have $\frac{\mathrm{d}}{\mathrm{d}t^0}(n_\perp \boldsymbol{u}_\perp^1) = \partial_{t^0}(n_\perp \boldsymbol{u}_\perp^1) = 0$, by invoking (3.3.7) and the fact that $\boldsymbol{u}_\perp^1$ is expressed in terms of the Maxwellian moments. The latter three terms of (3.4.15) all vanish by the same argument. Therefore $\hat{f}^2$ has strictly even-order polynomial dependence on $\boldsymbol{v}_\perp$, and the final term of (3.4.20) vanishes, so that $\boldsymbol{q}_\perp^2 = 0$ in the case of parallel symmetry and negligible collisions.

3.5 Summary of semi-fluid systems

3.5.1 Diffusive expansion

The diffusive expansion results in a pair of equations for the density and energy:

$$\frac{\mathrm{d}}{\mathrm{d}t}n_\perp + \nabla_\perp \cdot (n_\perp \boldsymbol{u}_\perp^1) = \mathcal{N}^0 \tag{3.5.1}$$

$$\frac{\mathrm{d}}{\mathrm{d}t}e_\perp = n_\perp Z\boldsymbol{E}\cdot\boldsymbol{u}_\perp^1 + \mathcal{Q}^0. \tag{3.5.2}$$

The first-order drift velocity is defined as

$$\boldsymbol{u}_\perp^1 = \frac{\boldsymbol{B}\times p_\perp^0}{n_\perp \omega_c \tau Z|B|^2} + \frac{\boldsymbol{E}\times\boldsymbol{B}}{\omega_c\tau|B|^2}, \tag{3.5.3}$$

where $p_\perp^0 = e_\perp$. Recall that $\boldsymbol{u}_\perp^0 = 0$ by our assumption on the leading-order distribution function. The collisional source terms $\mathcal{N}^0$ and $\mathcal{Q}^0$ are simply collisional moments of the gyrotropic Maxwellian. They may be calculated by direct integration for any given collision operator. For a bilinear collision operator C, for example, we have

$$\mathcal{N}_a^0 = \int C(f_a^0, f_b^0)\,\mathrm{d}\boldsymbol{v}_\perp, \quad \mathcal{Q}_a^0 = \frac{A}{2}\int |\boldsymbol{v}_\perp|^2 C(f_a^0, f_b^0)\,\mathrm{d}\boldsymbol{v}_\perp.$$

In the absence of collisions and parallel variation, i.e. when $\mathcal{N}^0 = \mathcal{Q}^0 = 0$ and $\frac{\mathrm{d}}{\mathrm{d}t} = \partial_t$, it is notable that the continuity equation is completely decoupled from the energy equation, since the divergence of the diamagnetic momentum vanishes. In this regime it is valid to solve (3.5.1) alone, as a very simple "drift-advection" model in the perpendicular direction.

3.5.2 Hyperbolic expansion

The hyperbolic expansion gives the usual five-moment system of equations for the perpendicular moments:

$$\frac{\mathrm{d}}{\mathrm{d}t}n_\perp + \nabla_\perp \cdot (n_\perp \boldsymbol{u}_\perp) = \mathcal{N} \tag{3.5.4}$$

$$A\frac{\mathrm{d}}{\mathrm{d}t}(n_\perp\boldsymbol{u}_\perp) + \nabla_\perp\cdot(An_\perp\boldsymbol{u}_\perp\otimes\boldsymbol{u}_\perp + \mathbb{P}_\perp) = n_\perp Z(\boldsymbol{E} + \omega_c\tau\boldsymbol{u}_\perp\times\boldsymbol{B}) + \mathcal{S} \tag{3.5.5}$$

$$\frac{\mathrm{d}}{\mathrm{d}t}e_\perp + \nabla_\perp\cdot((e_\perp\mathbb{I} + \mathbb{P}_\perp)\cdot\boldsymbol{u}_\perp + \boldsymbol{q}_\perp) = n_\perp Z\boldsymbol{E}\cdot\boldsymbol{u}_\perp + \mathcal{Q}. \tag{3.5.6}$$

The heat flux $\boldsymbol{q}_\perp$ is given by

$$\boldsymbol{q}_\perp = 2p_\perp^0\frac{\boldsymbol{B}\times\nabla_\perp T_\perp^0}{\omega_c\tau Z|B|^2} + \boldsymbol{q}_\perp^2, \tag{3.5.7}$$

where $p_\perp^0 = n_\perp T_\perp^0 = e_\perp$. The second-order heat flux $\boldsymbol{q}_\perp^2$ vanishes in the absence of collisions and parallel dynamics.

The pressure tensor is given by $\mathbb{P}_\perp = p_\perp \mathbb{I} + \Pi_\perp$, where $p_\perp = e_\perp - \frac{An_\perp |\boldsymbol{u}_\perp|^2}{2}$, and

$$\Pi_\perp = p_\perp^0 \frac{\mathbb{W}_3[\boldsymbol{u}_\perp]}{2\Omega_c} + \frac{\mathbb{W}_3[\boldsymbol{q}_\perp]}{4\Omega_c} + \frac{An_\perp}{2}[\widehat{(\boldsymbol{u}_\perp^\dagger - \boldsymbol{u}_\perp) \otimes \boldsymbol{u}_\perp}] + \frac{A}{2p_\perp^0 T_\perp^0}(\widehat{\boldsymbol{q}_\perp \otimes \boldsymbol{q}_\perp}), \tag{3.5.8}$$

where $\boldsymbol{u}_\perp^\dagger = \boldsymbol{u}_\perp^1 + \boldsymbol{u}_T$ and $\Omega_c = Z\omega_c\tau/A$ is the dimensionless species cyclotron frequency. The first-order drift velocity $\boldsymbol{u}_\perp^1$ is defined as in the diffusive expansion via (3.5.3).

It is interesting to compare the closures (3.5.7) and (3.5.8) with Braginskii's closures. The diamagnetic heat flux closure matches Braginskii's result after adjusting for the number of degrees of freedom, with the exception that the total energy quantities $p_\perp^0$ and $T_\perp^0$ are used instead of internal energy quantities. This amounts to an order ϵ^2 discrepancy in practice if the velocity $\boldsymbol{u}_\perp$ is small.

The pressure tensor closure (3.5.8) is different: the first term corresponds to Braginskii's gyroviscosity closure, which is the η_3 term of his Equation (4.41) [11]. However, the second, third, and fourth terms of (3.5.8) have no analogue in Braginskii's expansion. The second term indicates a contribution to the gyroviscosity attributable to a temperature gradient. A simple qualitative explanation of the third and fourth terms seems difficult to provide.

As with the diffusive scaling, the collisional source terms $\mathcal{N}, \mathcal{S}, \mathcal{Q}$ can be calculated as integrals of the collision operator. We do not elaborate on collisional transport here.

3.6 Discussion

We have derived a pair of asymptotic closures for the perpendicular moments of a magnetized, collisionless plasma. The first, which we term the diffusive expansion, leads to a drift-advection equation for particle density which, in the absence of collisions, is decoupled from the energy equation and may be solved independently. The second is a five-moment system for the perpendicular moments including a momentum equation. The gyroviscous stress is found to differ from that predicted by Braginskii. In this section we compare the FLR semi-fluid transport theory derived in this chapter to other reduced models of kinetic physics in the strongly magnetized regime.

Relation to gyrokinetic theory

Gyrokinetics is an extremely successful model reduction for strongly magnetized plasmas. The essence of the gyrokinetic approach is to average over the gyrophase of the magnetized particles and derive a kinetic equation for the particle guiding centers [28]. Gyrokinetics is widely used in tokamak modeling due to the enormous scale separation between device time scales and ion (let alone electron) cyclotron frequencies. As in the rest of this chapter, we neglect effects resulting from non-straight magnetic field lines.

We begin by noting that the semi-fluid point of view, which retains the kinetic description for the parallel velocity direction, has much in common with a gyrokinetic description. For example, the semi-fluid continuity equation

$$\partial_t n_\perp + \nabla_\perp \cdot (n_\perp \boldsymbol{u}_\perp) + v_\parallel \nabla_\parallel n_\perp = \mathcal{N}^0$$

can be derived by taking the zeroth μ moment of a gyrokinetic equation such as

$$\partial_t F + \mathbf{v}_\perp \cdot \nabla_{\boldsymbol{R}} F + v_\parallel \nabla_\parallel F = \tilde{C},$$

where $F(\boldsymbol{R}, \mu, v_\parallel, t)$ is the guiding center distribution, $\boldsymbol{R}$ is the guiding center spatial position, and μ is the magnetic moment. Higher-order moments of μ have a similar correspondence to *even-order* moments of $\boldsymbol{v}_\perp$, such as the semi-fluid pressure equation which could be obtained by manipulating (3.2.10) and (3.2.11).

However, gyrokinetics does not include a self-consistent equation for the perpendicular momentum of guiding centers. The gyrokinetic change of coordinates to particle guiding centers necessarily imports an asymptotic theory of how those guiding centers behave. For example, at leading order in Larmor radius, it is correct to model guiding centers as having velocities coming from the $E \times B$, ∇B, and magnetic curvature drifts. Higher-order effects such as the ion polarization drift must be included explicitly as an additional term. However, it is possible to perform similar transport analyses from a gyrokinetic starting point; for example, Belova [3] derived the gyroviscous tensor and force for a gyrokinetic plasma.

In this respect, gyrokinetic theory is similar to the diffusive expansion of 3.3. Indeed, in the absence of collisions, the semi-fluid drift-advection equation (3.3.15) is equivalent to a long-wavelength gyrokinetic (sometimes called a drift-kinetic) description for uniform magnetic field. The hyperbolic expansion of 3.4, on the other hand, includes a self-consistent treatment of momentum.

A drawback of the semi-fluid approach as compared to gyrokinetics is that it only describes up to the second moment of perpendicular velocity. By including the magnetic moment μ as a coordinate, gyrokinetic models can describe arbitrarily non-Maxwellian, but gyrotropic, distribution functions. Our semi-fluid approach, on the other hand, is limited to distribution functions that are Maxwellian with an order ϵ departure from gyrotropy.

Based on this summary, we hazard a prescription:

- Gyrokinetic descriptions are far more suitable for describing the fine structure of the distribution function, particularly with regard to collisional physics. For example, the phenomenon of particle detrapping in mirrors due to collisions requires representing a loss cone distribution which is highly non-Maxwellian and cannot be easily represented with semi-fluid variables.
- The semi-fluid description may be more effective than gyrokinetics at capturing finite-Larmor radius

effects such as gyroviscosity and second-order and higher drift velocities. Moreover, the perpendicular moment-based approach reduces the number of degrees of freedom compared to treating μ as a coordinate.

Relation to gyrofluid models

Gyrofluid models of plasma have been developed by Hammett, Perkins and Dorland [55, 54, 28], with the goal of obtaining a moment-based description for the *parallel* evolution of plasma in highly magnetized regimes. Rather than take a Maxwellian distribution in $v_\parallel$ as an ansatz, gyrofluid closures choose closure moments to approximately replicate the electrostatic dispersion relation of a Maxwellian distribution.

The general idea of obtaining fluid closures for parallel moments is, we believe, equally applicable to gyrokinetic and semi-fluid equations. The resulting moment system will retain the same perpendicular moments as the underlying "kinetic" theory, be it gyrokinetic or semi-fluid. The specific gyrofluid theory developed in [28] retains only the zeroth and gyrotropic second moment of the guiding center distribution, and so neglects self-consistent perpendicular momentum in the same way as gyrokinetic theory. Developing a gyrofluid closure for the parallel velocity moments of our semi-fluid equations is outside the scope of this work.

Assumption of Maxwellian equilibrium

Both the diffusive expansion and the hyperbolic expansion make the assumption that the leading-order solution is a Maxwellian. This should be regarded as a modeling assumption which dictates the validity of the resulting theory. In this respect the Maxwellian assumption is similar to assuming the relative smallness of various terms in the kinetic equation, which is what we do when we introduce the parameter ϵ. If the flow can be regarded as only weakly compressible, for example, then assuming a nearly Maxwellian velocity distribution is more likely to be valid.

It may be possible to derive a transport theory by assuming that the zeroth-order solution is arbitrarily gyrotropic, i.e. by assuming an expansion in terms of Laguerre polynomials:

$$f^0(\boldsymbol{x}, e_\perp, t) = e^{-e_\perp}\left[c_0(\boldsymbol{x},t)L_0^{(1/2)}(e_\perp) + c_1(\boldsymbol{x},t)L_1^{(1/2)}(e_\perp) + c_2(\boldsymbol{x},t)L_2^{(1/2)}(e_\perp) + \cdots\right].$$

However, such an expansion would necessarily involve hard-to-interpret quantities such as c_2, which is related to the fourth moment of velocity. The Maxwellian ansatz strikes a balance between physical plausibility and interpretability which leads to a highly tractable set of fluid equations.

Conclusion

The semi-fluid equations derived in this chapter serve as alternatives to the often-used Braginskii transport equations. Insofar as we have not assumed anything about the strength of collisions to isotropize the distribution function between parallel and perpendicular velocities, the assumptions made here are less restrictive than the assumptions made by the Braginskii fast-dynamics ordering. On the other hand, the resulting equations are kinetic in parallel velocity, and so in that sense less progress has been made to reduce the kinetic model to a fluid model. The semi-fluid equations are therefore particularly suited to modeling particle, energy and momentum transport in the perpendicular direction, such as when the parallel direction is assumed to be a direction of symmetry. Such applications may include modeling of tokamaks and magnetically-insulated transmission lines [108].

Further work is required to extend the semi-fluid closure model to more magnetically dynamic situations. In particular, magnetic curvature and changing magnetic fields introduce coordinate transform terms when separating $\boldsymbol{v}_\perp$ from $v_\parallel$ which complicate the derivation. We are optimistic about the possibility of generalizing these calculations, particularly if the magnetic field gradients and time derivatives can be ordered small relative to the cyclotron frequency and Larmor radius.

Another extension of the work presented here relates to calculating collisional moments. We have left the collisional moments $\mathcal{N}, \mathcal{S}$, and $\mathcal{Q}$ abstract, in the sense of not substituting any particular collision operator and taking its moments. However, collisional physics plays a key role in viscosity and resistivity, even in plasma regimes where collisions are not strong enough to send f to a Maxwellian. A key question is how the collisional transport contained in this semi-fluid model compares to the collisional transport terms in the Braginskii closure.

Chapter 4

A Fourier-Hermite spectral Vlasov solver

This chapter describes a simulation framework for the Vlasov-Fokker-Planck equation in the electrostatic approximation. The numerical discretization and implementation described here have been designed with an eye towards simulating magnetized plasmas at arbitrary collisionalities.

Our focus is on the so-called "slab" geometry, by which we mean a box in Cartesian geometry, with a straight magnetic field aligned to one of the coordinate axes. The Cartesian unit vectors $(\hat{\boldsymbol{x}}, \hat{\boldsymbol{y}}, \hat{\boldsymbol{z}})$ form the usual right-handed coordinate system, and we take $\boldsymbol{B}$ to be oriented in the $\hat{\boldsymbol{y}}$ direction, i.e. $\boldsymbol{B} = B\hat{\boldsymbol{y}}$. The $\hat{\boldsymbol{x}}$ and $\hat{\boldsymbol{y}}$ dimensions are both taken to be periodic, while the $\hat{\boldsymbol{z}}$ dimension is finite and able to assume boundary conditions for the $\boldsymbol{E}$ field and for the distribution functions $f_{e,i}$.

The choice of a rectangular geometry allows us to use a structured grid in $\boldsymbol{x}$. Our discretization of phase space is dictated by the choice of a structured grid and by the types of problems we consider. In particular, the problems to which we apply this code do not involve strong shocks or discontinuities in physical space. For the periodic dimensions, we therefore choose a Fourier pseudospectral discretization, while the lone non-periodic spatial dimension is discretized with a simple high-order conservative finite difference scheme. Meanwhile, in velocity space, we do not seek to resolve strong phase-space mixing or filamentation. We therefore opt for a Hermite spectral discretization of velocity. Such a discretization is simple to implement, highly accurate for the types of problems we consider, and efficient on GPU architectures.

For the problems we consider, the "dynamic" time τ_{dyn}, i.e. the timescale at which the phenomena of interest occur, is within two or three orders of magnitude of the inverse *ion* cyclotron frequency. However, a full Vlasov discretization for the electron species requires resolving the electron cyclotron frequency, Because of the very small mass ratio m_e/m_i between electrons and ions, the inverse electron cyclotron frequency is four or five orders of magnitude separated from τ_{dyn}. This represents a very substantial performance hit for little benefit in terms of additional physics of interest. To avoid the need to resolve these fast electron timescales, we implement support for simplified "drift-advection" modeling of electrons via Equation (3.3.15).

For convenience, we repeat the normalized Vlasov equation

$$\partial_t f_a + \boldsymbol{v} \cdot \nabla_{\boldsymbol{x}} f_a + \frac{Z_a}{A_a} \left[(\omega_p \tau) \boldsymbol{E} + (\omega_c \tau) \boldsymbol{v} \times \boldsymbol{B} \right] \cdot \nabla_{\boldsymbol{v}} f_a = 0, \tag{4.0.1}$$

and the collisionless drift-advection equation

$$\partial_t n_a + \nabla \cdot (n_a \boldsymbol{u}_E) = 0, \tag{4.0.2}$$

where

$$\boldsymbol{u}_E = \frac{\omega_p \tau \boldsymbol{E} \times \boldsymbol{B}}{\omega_c \tau |B|^2}$$

is the $E \times B$ drift velocity. We do not consider problems involving a magnetic field gradient or curvature, and so neglect the resulting drifts.

4.1 Physical space discretization

Our discretization of physical space is designed to leverage the substantial advantages afforded by a rectangular "slab" geometry. Since we do not need to handle complex boundary shapes, we are free to use a structured grid. We therefore use a simple and high-order accurate tensor-product discretization, which combines a Fourier pseudospectral discretization of the periodic x and y dimensions with a conservative finite difference discretization of the non-periodic z dimension. An excellent introduction to pseudospectral discretizations in the Fourier basis can be found in [10]. Our conservative finite difference discretization is in what is known as Shu-Osher form [93]. For an interesting discussion of the merits of the Shu-Osher finite difference scheme, in particular its surprising efficiency in multiple dimensions, we refer the reader to [82].

To illustrate the physical space discretization, consider a vector-valued model equation for an unknown $\boldsymbol{q}$, with a linear flux term and source term:

$$\partial_t \boldsymbol{q}(\boldsymbol{x}, t) + \nabla \cdot (\boldsymbol{V} \boldsymbol{q}) = \boldsymbol{S}(\boldsymbol{q}). \tag{4.1.1}$$

The drift-advection equation (4.0.2) is evidently in this form. Such an equation also results from a velocity space discretization of the normalized Vlasov-Boltzmann (4.0.1) as we will see in Section 4.2. The matrix $\boldsymbol{V}$ is the flux Jacobian matrix, and represents the free-streaming or "multiplication-by-v" operator in the case of the Vlasov equation, and multiplication by the drift velocity in the drift-advection case. The divergence of the flux can be expanded as follows:

$$\partial_t \boldsymbol{q} + \partial_x (V_x \boldsymbol{q}) + \partial_y (V_y \boldsymbol{q}) + \partial_z (V_z \boldsymbol{q}) = \boldsymbol{S}(\boldsymbol{q}), \tag{4.1.2}$$

where V_x, V_y, V_z are the flux Jacobians in each of the spatial dimensions. Let the spatial domain of $\boldsymbol{q}(x, y, z)$ be

$$(x, y, z) \in \Omega = \Omega_x \times \Omega_y \times \Omega_z,$$

where $\Omega_x = [0, L_x)$ is the periodic interval with length L_x and likewise $\Omega_y = [0, L_y)$ the periodic interval with length L_y. We begin by discretizing the periodic dimensions x and y, leaving z continuous for the time being. Our Fourier pseudospectral discretization approximates $\boldsymbol{q}$ as a tensor product of trigonometric polynomials in x and y, whose coefficients are continuous functions of z:

$$\begin{align} \boldsymbol{q}_h(x, y, z) &= \sum_{k_x=-N_x/2}^{N_x/2-1} \sum_{k_y=-N_y/2}^{N_y/2-1} e^{\frac{2\pi i k_x x}{L_x}} e^{\frac{2\pi i k_y y}{L_y}} \hat{\boldsymbol{q}}_{k_x k_y}(z) \tag{4.1.3} \\ &\triangleq \mathcal{F}^{-1}[\hat{\boldsymbol{q}}_{k_x k_y}(z)](x, y). \tag{4.1.4} \end{align}$$

The subscript h denotes that $\boldsymbol{q}_h$ is an approximate solution which depends on the grid spacing. The Fourier modes $\hat{\boldsymbol{q}}_{k_x k_y}(z)$ take complex values. They are related to the point values of $\boldsymbol{q}_h$ at an evenly spaced set of grid points (x_i, y_j) via the discrete Fourier transform (DFT). We have already defined the inverse DFT via (4.1.4). The forward DFT maps function values $\boldsymbol{q}_h(x_i, y_j, z)$ to Fourier modes:

$$\begin{aligned} \hat{\boldsymbol{q}}_{k_x,k_y}(z) &= \mathcal{F}[\boldsymbol{q}(x_i, y_j, z)](k_x, k_y) \\ &= \frac{1}{(2\pi)^2 N_x N_y} \sum_{i=0}^{N_x-1} \sum_{j=0}^{N_y-1} e^{-\frac{2\pi i k_x x_i}{L_x}} e^{-\frac{2\pi i k_y y_j}{L_y}} \boldsymbol{q}_h(x_i, y_j, z). \end{aligned} \tag{4.1.5}$$

With any spectral discretization, the second choice to be made (after a choice of function space) is whether the discrete solution ought to satisfy the governing equation in a Galerkin or a collocation sense. For a Galerkin spectral scheme, one requires the residual of the governing equation to be orthogonal to the approximation space. A collocation method requires the governing equation to be satisfied at a discrete set of points. For a discussion of the merits of both approaches, [10] is a good reference. In this work we opt for a collocation scheme. This is in part for harmony with our discretization of the z dimension, which is a finite difference scheme. The overall scheme is greatly simplified by treating each degree of freedom in x, y, z as a function value.

We have already introduced the grid points x_i and y_j in the context of the discrete Fourier transform. The z dimension is also equipped with an evenly spaced set of grid points, z_k. We therefore ask that (4.1.3) satisfy

(4.1.2) in a collocation, or pointwise, sense at the evenly spaced grid points (x_i, y_j, z_k). As we will see, the actual degrees of freedom in our code are the function values, $\boldsymbol{q}_h(x_i, y_j, z_k)$, which can be converted back and forth to the Fourier modes $\hat{\boldsymbol{q}}_{ij}(z_k)$ efficiently using the Discrete Fourier Transform, defined by (4.1.4) and (4.1.5).

The strategy for the pseudospectral collocation scheme is as follows. To evaluate operators on the grid function $\boldsymbol{q}_h$, we perform the evaluation pointwise at the collocation points (x_i, y_j, z_k). The x and y flux derivatives are evaluated in the Fourier domain. Thus, we have

$$\begin{aligned} \partial_x(V_x\boldsymbol{q}_h)(x_i, y_j, z) &= \mathcal{F}^{-1}\left[\frac{2\pi i k_x}{L_x}\mathcal{F}\left\{V_x\boldsymbol{q}_h\right\}(k_x, k_y, z)\right](x_i, y_j, z), \\ \partial_y(V_y\boldsymbol{q}_h)(x_i, y_j, z) &= \mathcal{F}^{-1}\left[\frac{2\pi i k_y}{L_y}\mathcal{F}\left\{V_y\boldsymbol{q}_h\right\}(k_x, k_y, z)\right](x_i, y_j, z). \end{aligned} \tag{4.1.6}$$

In the case of the Vlasov equation, the operator matrices V_x and V_y are independent of $\boldsymbol{x}$, so we are free to pull the application of V_x and V_y outside the Fourier transforms. This is not the case for the drift-kinetic formulation however, so the conservative discretization (4.1.6) is the more general formulation.

The z flux derivative is evaluated using a fifth-order Shu-Osher conservative finite difference method [93, 92].

$$\partial_z F(\boldsymbol{q}_h)|_{z=z_k} = \frac{1}{\Delta z}\left(\hat{F}_{k+1/2} - \hat{F}_{k-1/2}\right).$$

The numerical flux $\hat{F}$ is split into a left-going and right-going part, which are reconstructed from upwind-biased stencils:

$$\hat{F}_{k+1/2} = \hat{F}^{+}_{k+1/2} + \hat{F}^{-}_{k+1/2},$$

where

$$\hat{F}^{+}_{k+1/2} = \frac{F^{+}_{k-2}}{30} - \frac{13F^{+}_{k-1}}{60} + \frac{47F^{+}_{k}}{60} + \frac{9F^{+}_{k+1}}{20} - \frac{F^{+}_{k+2}}{20}$$

and

$$\hat{F}^{-}_{k+1/2} = -\frac{F^{-}_{k-1}}{20} + \frac{9F^{-}_{k}}{20} + \frac{47F^{-}_{k+1}}{60} - \frac{13F^{-}_{k+2}}{60} + \frac{F^{-}_{k+3}}{30}.$$

The splitting of the analytic flux $F(\boldsymbol{q}_h) = F^{-} + F^{+}$ can be accomplished in various ways. A natural choice for a linear flux such as $F(\boldsymbol{q}_h) = V_z\boldsymbol{q}_h$ is to perform direct characteristic decomposition of V_z, such that

$$F^{\pm}(\boldsymbol{q}_h) = R\Lambda^{\pm}R^{-1},$$

where $\Lambda^{\pm}$ is given by

$$\Lambda^{\pm} = \frac{1}{2}\left(\Lambda \pm |\Lambda|\right),$$

with $V_z = R\Lambda R^{-1}$ the eigendecomposition of the matrix. This choice is stable and introduces little extra numerical diffusion.

However, for the Hermite spectral velocity discretization we consider here, this choice destroys the sparsity of the flux matrix: while $R\Lambda R^{-1}$ is tridiagonal, $R\Lambda^{\pm}R^{-1}$ is a dense matrix. Therefore, for efficiency it is better to choose a Lax-Friedrichs flux:

$$F^{\pm}(\boldsymbol{q}_h) = \frac{1}{2}\left(V_z \pm \alpha I\right),$$

where $\alpha = \max_{\lambda \in \mathrm{diag}(\Lambda)} |\lambda|$ is the maximum absolute eigenvalue of V_z. With the Lax-Friedrichs flux, the upwind and downwind fluxes $F^{\pm}$ can still be calculated using a tridiagonal matrix.

4.2 Velocity space discretization of Vlasov species

There are many viable choices for numerical discretization of velocity space for the Vlasov and Vlasov-Boltzmann equations. Because the acceleration term $\frac{A}{A}(\omega_p \tau \boldsymbol{E} + \omega_c \tau \boldsymbol{v} \times \boldsymbol{B}) \cdot \nabla_{\boldsymbol{v}} f$ is only weakly nonlinear in f, the nonlinear coupling occuring through the $\boldsymbol{E}$ and $\boldsymbol{B}$ fields, relatively simple numerical methods can give good accuracy.

However, it is natural to insist on a method that is exactly or nearly exactly conservative of the first three moments of the distribution function,

$$M_0 = \int f \,\mathrm{d}\boldsymbol{v}, \quad \boldsymbol{M}_1 = \int \boldsymbol{v} f \,\mathrm{d}\boldsymbol{v}, \quad M_2 = \int \frac{|\boldsymbol{v}|^2}{2} f \,\mathrm{d}\boldsymbol{v}, \tag{4.2.1}$$

which are proportional to the mass density, local momentum, and kinetic energy density of the distribution function f. Conserving *functionals* of a PDE solution such as (4.2.1) is more challenging than coming up with a locally conservative discretization of f itself.

One family of methods that is capable of conserving the functionals defined in (4.2.1) is the Discontinuous Galerkin (DG) method [18, 58]. The DG method for the Vlasov equation has been implemented in the WARPXM code [59, 24] and the Gkeyll code [68, 51]. A conservative DG discretization of the Lenard-Bernstein operator is presented in [52]. The advantages of Discontinuous Galerkin discretizations include arbitrarily high-order accuracy, a good ability to capture fine features in velocity space, and the aforementioned conservation properties. However, for the problems that we consider, Runge-Kutta DG (RKDG) methods suffer from arduous stability restrictions for the explicit timestep.

In this work, rather than a grid-based velocity space discretization such as DG, finite difference or finite volume, we make use of a global Hermite spectral discretization of velocity space. Hermite polynomial representations have been used before to discretize the Vlasov equation, including methods that couple a Hermite velocity space representation with DG [74, 39] and Fourier spectral [105, 26] representations in physical space. Our discretization is distinct in that it combines a Hermite Galerkin spectral scheme in velocity space, a Fourier collocation spectral scheme in the x and y space dimensions, and a high-order finite difference scheme in z. The choice of a Galerkin projection in velocity space leads to a concise evaluation of the upwind velocity flux in the z direction via a Lax-Friedrichs flux splitting.

The discretization expands a function $f(\boldsymbol{x}, v_x, v_y, v_z)$ in terms of a tensor product of asymmetrically-weighted normalized Hermite polynomials:

$$f(\boldsymbol{x}, v_x, v_y, v_z) = \sum_{l=0}^{N_{v_x}} \sum_{m=0}^{N_{v_y}} \sum_{n=0}^{N_{v_z}} \boldsymbol{f}_{lmn}(\boldsymbol{x}) \frac{\exp\left(-\frac{(v_x^2+v_y^2+v_z^2)}{2v_{th}^2}\right)}{(2\pi)^{3/2} v_{th}^3} He_l\left(\frac{v_x}{v_{th}}\right) He_m\left(\frac{v_y}{v_{th}}\right) He_n\left(\frac{v_z}{v_{th}}\right), \tag{4.2.2}$$

where v_{th} is a reference thermal velocity which sets the width of the basis functions in velocity space. The first several normalized Hermite polynomials are

$$\begin{aligned} He_0\left(\frac{v}{v_{th}}\right) &= 1 & He_2\left(\frac{v}{v_{th}}\right) &= \frac{(v/v_{th})^2 - 1}{\sqrt{2}} \\ He_1\left(\frac{v}{v_{th}}\right) &= v/v_{th} & He_3\left(\frac{v}{v_{th}}\right) &= \frac{(v/v_{th})^3 - 3v/v_{th}}{\sqrt{6}}. \end{aligned}$$

As a family of polynomials, they share the property that they are orthogonal with respect to the weight function $\frac{1}{\sqrt{2\pi} v_{th}} e^{-\frac{v^2}{2v_{th}^2}}$:

$$\frac{1}{\sqrt{2\pi} v_{th}} \int_{\mathbb{R}} He_i\left(\frac{v}{v_{th}}\right) He_j\left(\frac{v}{v_{th}}\right) e^{-\frac{v^2}{2v_{th}^2}} \,\mathrm{d}v = \delta_{ij}. \tag{4.2.3}$$

The orthogonality relation (4.2.3) implies that the degrees of freedom $\boldsymbol{f}_{lmn}(\boldsymbol{x})$ are not only coefficients of the Hermite expansion, but also *moments* of f with respect to the Hermite polynomials. Indeed,

$$\begin{aligned} &\int_{\mathbb{R}^3} He_p\left(\frac{v_x}{v_{th}}\right) He_q\left(\frac{v_y}{v_{th}}\right) He_r\left(\frac{v_z}{v_{th}}\right) f(\boldsymbol{x}, v_x, v_y, v_z) \,\mathrm{d}\boldsymbol{v} \\ &= \sum_{l=0}^{N_{v_x}} \sum_{m=0}^{N_{v_y}} \sum_{n=0}^{N_{v_z}} \boldsymbol{f}_{lmn}(\boldsymbol{x}) \left[\int_{\mathbb{R}} He_p\left(\frac{v_x}{v_{th}}\right) He_l\left(\frac{v_x}{v_{th}}\right) \frac{e^{-\frac{v_x^2}{2v_{th}^2}}}{\sqrt{2\pi} v_{th}} \,\mathrm{d}v_x \right] \\ &\quad \times \left[\int_{\mathbb{R}} He_q\left(\frac{v_y}{v_{th}}\right) He_m\left(\frac{v_y}{v_{th}}\right) \frac{e^{-\frac{v_y^2}{2v_{th}^2}}}{\sqrt{2\pi} v_{th}} \,\mathrm{d}v_y \right] \left[\int_{\mathbb{R}} He_r\left(\frac{v_z}{v_{th}}\right) He_n\left(\frac{v_z}{v_{th}}\right) \frac{e^{-\frac{v_z^2}{2v_{th}^2}}}{\sqrt{2\pi} v_{th}} \,\mathrm{d}v_z \right] \\ &= \sum_{l=0}^{N_{v_x}} \sum_{m=0}^{N_{v_y}} \sum_{n=0}^{N_{v_z}} \boldsymbol{f}_{lmn}(\boldsymbol{x}) \delta_{lp} \delta_{mq} \delta_{nr} \\ &= \boldsymbol{f}_{pqr}(\boldsymbol{x}). \end{aligned}$$

A consequence of this property is that a conservative Hermite spectral discretization—that is, a discretization which conserves the Hermite *coefficients* $\boldsymbol{f}_{lmn}(\boldsymbol{x})$—will automatically conserve the mass, momentum, and energy of the distribution function. For, by linearity of the moment-taking functional, we can deduce that

$$M_0 = \boldsymbol{f}_{000}, \quad \boldsymbol{M}_1 = v_{th}(\boldsymbol{f}_{100}, \boldsymbol{f}_{010}, \boldsymbol{f}_{001})^T, \quad M_2 = v_{th}^2(\sqrt{2}(\boldsymbol{f}_{200} + \boldsymbol{f}_{020} + \boldsymbol{f}_{002}) + \boldsymbol{f}_{000}).$$

In other words, the conserved functionals of f defined in (4.2.1) naturally appear in the Hermite discretization as linear combinations of the degrees of freedom. This property makes it easy to obtain a fully discrete scheme with good conservation properties.

We now describe the implementation of the free streaming and derivative operators in the Hermite spectral discretization, which are implemented as tensor products of one-dimensional operators. For a one-dimensional function $h(v)$ expressed as a sum of asymmetrically-weighted Hermite polynomials,

$$h(v) = \frac{1}{\sqrt{2\pi} v_{th}} \sum_i h_i(v) He_i\left(\frac{v}{v_{th}}\right) e^{-\frac{v^2}{2v_{th}^2}},$$

the "multiplication-by-v" operator can be expressed as a matrix-vector product acting on the vector of coefficients:

$$vh(v) = \frac{1}{\sqrt{2\pi} v_{th}} \sum_i \sum_j v_{th} V_{ij}^H h_j(v) He_i\left(\frac{v}{v_{th}}\right) e^{-\frac{v^2}{2v_{th}^2}}. \tag{4.2.4}$$

The entries of the matrix V can be computed from the recurrence relations of the Hermite polynomials [27]:

$$V^H = \begin{pmatrix} 0 & 1 & & & \\ 1 & 0 & \sqrt{2} & & \\ & \sqrt{2} & 0 & \sqrt{3} & \\ & & \sqrt{3} & 0 & \ddots \\ & & & \ddots & \ddots \end{pmatrix}. \tag{4.2.5}$$

The ∂_v operator can also be written as a matrix-vector product,

$$\partial_v h(v) = \frac{1}{\sqrt{2\pi} v_{th}} \sum_i \sum_j \frac{1}{v_{th}} D_{ij}^H h_j(v) He_i\left(\frac{v}{v_{th}}\right) e^{-\frac{v^2}{2v_{th}^2}}, \tag{4.2.6}$$

where

$$D^H = \begin{pmatrix} 0 & & & & \\ -1 & 0 & & & \\ & -\sqrt{2} & 0 & & \\ & & -\sqrt{3} & 0 & \\ & & & \ddots & \ddots \end{pmatrix}. \tag{4.2.7}$$

With these discrete operators, we are ready to write the semi-discrete form of the Vlasov-Boltzmann equation. To avoid overly cumbersome notation, we introduce the shorthand

$$\mathcal{H}_i^a(v) = \frac{1}{\sqrt{2\pi} v_{th,a}} e^{-\frac{v^2}{v_{th,a}^2}} He_i\left(\frac{v}{v_{th,a}}\right).$$

Note the reapparance of the species index a, which reflects the fact that the reference velocity $v_{th,a}$ may be chosen differently for different species. In terms of this notation, our ansatz for f_a is

$$f_a(\boldsymbol{x}, v_x, v_y, v_z) = \boldsymbol{f}^a_{lmn}(\boldsymbol{x}) \mathcal{H}^a_l(v_x) \mathcal{H}^a_m(v_y) \mathcal{H}^a_n(v_z), \tag{4.2.8}$$

where we are using the convention that repeated indices are summed over. Substituting (4.2.8) into (4.0.1), and making use of the relations (4.2.4) and (4.2.6), after some algebra we obtain

$$\begin{aligned}
\partial_t \boldsymbol{f}_{lmn} \mathcal{H}^a_l \mathcal{H}^a_m \mathcal{H}^a_n &+ \left(V^H_{lp} \partial_x \boldsymbol{f}^a_{pmn} + V^H_{mq} \partial_y \boldsymbol{f}^a_{lqn} + V^H_{nr} \partial_z \boldsymbol{f}^a_{lmr}\right) \mathcal{H}^a_l \mathcal{H}^a_m \mathcal{H}^a_n \\
&+ \frac{Z_a}{A_a} \frac{\omega_p \tau}{v_{th,a}} \left[E_x(\boldsymbol{x}) D^H_{lp} \boldsymbol{f}^a_{pmn} + E_y(\boldsymbol{x}) D^H_{mq} \boldsymbol{f}^a_{lqn} + E_z(\boldsymbol{x}) D^H_{nr} \boldsymbol{f}^a_{lmr}\right] \mathcal{H}^a_l \mathcal{H}^a_m \mathcal{H}^a_n \\
&+ \frac{Z_a}{A_a} (\omega_c \tau) \left[\left(B_z(\boldsymbol{x}) V^H_{mq} D^H_{lp} \boldsymbol{f}^a_{pqn} - B_y(\boldsymbol{x}) V^H_{nr} D^H_{lp} \boldsymbol{f}^a_{pmr}\right)\right] \mathcal{H}^a_l \mathcal{H}^a_m \mathcal{H}^a_n \\
&+ \frac{Z_a}{A_a} (\omega_c \tau) \left[\left(B_x(\boldsymbol{x}) V^H_{nr} D^H_{mq} \boldsymbol{f}^a_{lqr} - B_z(\boldsymbol{x}) V^H_{lp} D^H_{mq} \boldsymbol{f}^a_{pqn}\right)\right] \mathcal{H}^a_l \mathcal{H}^a_m \mathcal{H}^a_n \\
&+ \frac{Z_a}{A_a} (\omega_c \tau) \left[\left(B_y(\boldsymbol{x}) V^H_{lp} D^H_{nr} \boldsymbol{f}^a_{pmr} - B_x(\boldsymbol{x}) V^H_{mq} D^H_{nr} \boldsymbol{f}^a_{lqr}\right)\right] \mathcal{H}^a_l \mathcal{H}^a_m \mathcal{H}^a_n \\
&= (\nu_p \tau) \sum_b \nu_{ab} \left(\frac{T_{ab}}{m_a v_{th,a}^2} \left[\left(D^2\right)^H_{lp} \boldsymbol{f}^a_{pmn} + \left(D^2\right)^H_{mq} \boldsymbol{f}^a_{lqn} + \left(D^2\right)^H_{nr} \boldsymbol{f}^a_{lmr}\right] \right. \\
&\qquad + \left[(DV)^H_{lp} \boldsymbol{f}^a_{pmn} + (DV)^H_{mq} \boldsymbol{f}^a_{lqn} + (DV)^H_{nr} \boldsymbol{f}^a_{lmr}\right] \\
&\qquad \left. - \frac{1}{v_{th,a}} \left[u_{x,ab} D^H_{lp} \boldsymbol{f}^a_{pmn} + u_{y,ab} D^H_{mq} \boldsymbol{f}^a_{lqn} + u_{z,ab} D^H_{nr} \boldsymbol{f}^a_{lmr}\right] \right) \mathcal{H}^a_l \mathcal{H}^a_m \mathcal{H}^a_n.
\end{aligned} \tag{4.2.9}$$

Now, observing that every term in (4.2.9) is an expansion in the basis $\mathcal{H}^a_l \mathcal{H}^a_m \mathcal{H}^a_n$, we eliminate the continuous dependence on velocity by applying the operation

$$\star \mapsto \int_{\mathbb{R}^3} \star \, He_i\left(\frac{v_x}{v_{th}}\right) He_j\left(\frac{v_y}{v_{th}}\right) He_k\left(\frac{v_z}{v_{th}}\right) \mathrm{d}\boldsymbol{v}$$

to each term. After renaming the indices $(i, j, k) \mapsto (l, m, n)$ this results in a semidiscrete equation for $\boldsymbol{f}$, with continuous dependence on time and the physical space coordinate $\boldsymbol{x}$:

$$\begin{aligned}
\partial_t \boldsymbol{f}_{lmn}(\boldsymbol{x}, t) &+ \left(V^H_{lp} \partial_x \boldsymbol{f}^a_{pmn} + V^H_{mq} \partial_y \boldsymbol{f}^a_{lqn} + V^H_{nr} \partial_z \boldsymbol{f}^a_{lmr}\right) \\
&+ \frac{Z_a}{A_a} \frac{\omega_p \tau}{v_{th,a}} \left[E_x(\boldsymbol{x}) D^H_{lp} \boldsymbol{f}^a_{pmn} + E_y(\boldsymbol{x}) D^H_{mq} \boldsymbol{f}^a_{lqn} + E_z(\boldsymbol{x}) D^H_{nr} \boldsymbol{f}^a_{lmr}\right] \\
&+ \frac{Z_a}{A_a} (\omega_c \tau) \left[\left(B_z(\boldsymbol{x}) V^H_{mq} D^H_{lp} \boldsymbol{f}^a_{pqn} - B_y(\boldsymbol{x}) V^H_{nr} D^H_{lp} \boldsymbol{f}^a_{pmr}\right)\right] \\
&+ \frac{Z_a}{A_a} (\omega_c \tau) \left[\left(B_x(\boldsymbol{x}) V^H_{nr} D^H_{mq} \boldsymbol{f}^a_{lqr} - B_z(\boldsymbol{x}) V^H_{lp} D^H_{mq} \boldsymbol{f}^a_{pqn}\right)\right] \\
&+ \frac{Z_a}{A_a} (\omega_c \tau) \left[\left(B_y(\boldsymbol{x}) V^H_{lp} D^H_{nr} \boldsymbol{f}^a_{pmr} - B_x(\boldsymbol{x}) V^H_{mq} D^H_{nr} \boldsymbol{f}^a_{lqr}\right)\right] \\
&= (\nu_p \tau) \sum_b \nu_{ab} \left(\frac{T_{ab}}{m_a v^2_{th,a}} \left[\left(D^2\right)^H_{lp} \boldsymbol{f}^a_{pmn} + \left(D^2\right)^H_{mq} \boldsymbol{f}^a_{lqn} + \left(D^2\right)^H_{nr} \boldsymbol{f}^a_{lmr}\right] \right. \\
&\qquad + \left[(DV)^H_{lp} \boldsymbol{f}^a_{pmn} + (DV)^H_{mq} \boldsymbol{f}^a_{lqn} + (DV)^H_{nr} \boldsymbol{f}^a_{lmr}\right] \\
&\qquad \left. - \frac{1}{v_{th,a}} \left[u_{x,ab} D^H_{lp} \boldsymbol{f}^a_{pmn} + u_{y,ab} D^H_{mq} \boldsymbol{f}^a_{lqn} + u_{z,ab} D^H_{nr} \boldsymbol{f}^a_{lmr}\right] \right).
\end{aligned} \tag{4.2.10}$$

Examining the resulting semi-discrete system, we can see that it has the form of (4.1.1). We can therefore use the spatial discretization scheme of Section 4.1 to discretize in space and get a system of equations which is continuous only in time.

4.2.1 Advantages of Hermite spectral over grid-based discretizations

Like most explicit methods for hyperbolic equations, the maximum allowable timestep is proportional to the maximum wavespeed in the system. For a conservation law such as $\partial_t \boldsymbol{a} + \nabla \cdot \boldsymbol{F}(\boldsymbol{a}) = 0$, the

$$\Delta t_{max} \propto \frac{\lambda_{max}}{h} \frac{1}{2N + 1},$$

where h is the grid spacing, λ_{max} is the maximum absolute eigenvalue of the Jacobian $\boldsymbol{F}'(\boldsymbol{a})$ of the flux function, and N is the degree of polynomials used in the DG approximation [18].

For many of the problems we will consider, the most restrictive timestep comes from the cyclotron orbit of the ions. To complete one cyclotron orbit, a notional particle near the edge of the velocity distribution must traverse roughly $4N_v$ elements, and thus the maximum timestep is

$$\Delta t_{max} \approx \frac{1}{4 N_v \omega_c (2N + 1)}.$$

On the other hand, for a Hermite spectral expansion, the eigenvalues of both V and D scale with $\sqrt{N_v}$, so the the maximum allowable timestep is

$$\Delta t_{max} \approx \frac{1}{N_v \omega_c}.$$

Moreover, since the Hermite expansion is uniquely suited to resolving the Maxwellian distribution centered at the origin, we can use a resolution N_v significantly lower than a grid-based method such as DG. This matches the experience reported in e.g. [74], which uses only 9 Hermite modes to resolve distribution functions in benchmark problems such as the Orszag-Tang vortex.

4.3 Time discretization

To discretize (4.2.10) in time, we use an explicit third-order four stage Strong-Stability-Preserving [45] Runge-Kutta scheme [31]. For an ordinary autonomous differential equation $u'(t) = f(u)$, the scheme is defined as follows:

$$\begin{aligned} u^1 &= u^n + \frac{\Delta t}{2} f(u^n) && (4.3.1)\\ u^2 &= u^1 + \frac{\Delta t}{2} f(u^1) && (4.3.2)\\ u^3 &= \frac{2}{3}u^1 + \frac{1}{3}\left[u^2 + \frac{\Delta t}{2} f(u^2)\right] && (4.3.3)\\ u^{n+1} &= u^3 + \frac{\Delta t}{2} f(u^3). && (4.3.4) \end{aligned}$$

4.4 Filtering and hyperdiffusion

Following [39], we employ the filter known as Hou-Li's filter [60] which prescribes multiplying the m^{th} Hermite mode by a scaling factor $\sigma\left(\frac{m}{M+1}\right)$, where

$$\sigma(s) = \begin{cases} 1, & 0 \le s \le 2/3, \\ e^{-\beta s^\beta} & s > 2/3, \end{cases} \tag{4.4.1}$$

with $\beta = 36$ designed to eliminate the final mode to within machine precision. The Hou-Li filter is applied at every Runge-Kutta stage. Thus, in (4.3.1), it is applied to each of u^1, u^2, u^3 and u^{n+1}.

In addition to filtering the Fourier and Hermite spectral modes, we apply a hyperdiffusion operator in space. The hyperdiffusion operator modifies the right-hand side of (4.1.2) to

$$\partial_t \boldsymbol{q} + \partial_x(V_x \boldsymbol{q}) + \partial_y(V_y \boldsymbol{q}) + \partial_z(V_z \boldsymbol{q}) = \boldsymbol{S}(\boldsymbol{q}) - \eta \Delta^2(\boldsymbol{q}), \tag{4.4.2}$$

where Δ is the Laplacian,

$$\Delta = \partial_{xx} + \partial_{yy} + \partial_{zz}.$$

Hyperdiffusion is a commonly employed solution for avoiding aliasing instability in spectral discretizations of nonlinear problems [22]. We choose the parameter η so that the eigenvalue corresponding to grid-scale waves, namely waves with wavelength $3\Delta x$, is a fixed value η_{grid}. The discretization therefore retains its consistency to the non-diffusive PDE as the discretization is refined. We find that a value of $\eta_{grid} = 0.25$ suffices to stabilize all the problems we consider here.

4.5 Device-agnostic Julia implementation

We implement the simulation code in the Julia programming language [7]. Julia's multiple dispatch capabilities allow us to program in an array-centric style. Critically, this style lets the whole software package be abstract over the choice of CPU / GPU. For GPU support we make heavy use of CUDA.jl [5, 6]. Finite difference stencils are implemented using convolution primitives from NNlib.jl. Hermite basis operators V and D are implemented with sparse tridiagonal matrix multiplications. We use the FFTW[1] library for Discrete Fourier Transforms (DFTs) on the CPU, and CuFFT for DFTs on NVIDIA GPUs.

These device-agnostic building blocks are combined into kernels that evaluate each of the terms in the Vlasov equation. Listing 4.1 shows one such kernel. Line 12 gives an example of the `@.` macro. This macro invokes Julia's broadcasting feature, which uses stream fusion to compile the right-hand side expression into a single loop with no intermediate arrays.

```
1  function electrostatic_x!(df, f, Ex, sim, species::Species{<:Hermite}, buffer,
       xgrid_fft_plans)
2      (; discretization, q, m) = species
3      Nx, Ny, Nz, Nvx, Nvy, Nvz = size(discretization)
4      (; omega_p_tau, omega_c_tau, By) = sim
5  
6      no_escape(buffer) do
7          vzf = alloc_array(Float64, buffer, size(f)...)
8          @timeit "mul by vz" mul_by_vz!(vzf, f, discretization)
9  
10         F = alloc_array(Float64, buffer, Nx, Ny, Nz, Nvx, Nvy, Nvz)
11 
12         @. F = q / m * (omega_p_tau * Ex * f - omega_c_tau * vzf * By)
13 
14         df = reshape(df, (:, Nvx*Nvy*Nvz))
15         F = reshape(F, (:, Nvx*Nvy*Nvz))
16         mul!(df, F, discretization.vdisc.Dvx’, -1 / discretization.vdisc.vth, 1.0)
```

---

[1] The Fastest Fourier Transform in the West

```
    end
end
```

Listing 4.1: Kernel for calculating the $(E_x + v_z B_y)\partial_{v_x} f$ term.

4.6 Benchmark problems

The reproduction of linear instabilities is a very powerful methodology for benchmarking plasma simulations. These test cases provide a highly integrated test case for the plasma code, since many elements must come together correctly to reproduce the predicted linear response. In this section, we run our Fourier-Hermite spectral code through two such benchmarks, a weak Landau damping problem and the Dory-Guest-Harris instability for loss-cone equilibria.

4.6.1 Weak Landau damping

We begin our benchmarking with several examples of so-called “weak”, or linear Landau damping. This refers to the regime where $k\lambda_D \sim 1$, and the perturbation from equilibrium is sufficiently small that a linear analysis is predictive. We set up a Maxwellian equilibrium with unit density and temperature and zero drift velocity,

$$f_0(v) = \frac{1}{\sqrt{2\pi}} e^{-\frac{v^2}{2}}.$$

This is perturbed with a sinusoidal perturbation of amplitude α, so that

$$f(x, v, 0) = \frac{1 + \alpha \sin(kx)}{\sqrt{2\pi}} e^{-\frac{v^2}{2}}.$$

The distribution function is discretized with on the interval $[0, 2\pi)$ with 96 grid points in x and 256 Hermite modes in v_x. We initialize perturbations with wavenumbers of $k = \{0.5, 2/3, 1.0\}$. The thermal speed, Debye length, and normalized plasma frequency are all set to unity in these simulations. Each simulation is run to $\omega_p t = 50.0$, although for the cases of $k = 2/3$ and $k = 1$, the electric energy damps below machine precision before the simulation ends, or begins to exhibit the recurrence phenomenon [13]. For this reason the electric energy traces reported in Figure 4.1a are truncated for theses cases. The timestep is set to $\Delta t = 7 \times 10^{-4}$.

The electric energy traces are processed to find peaks, whose location in time and height lets us estimate both a damping rate and real oscillation frequency of the Langmuir wave. At the same time, the linearized Vlasov dispersion relation for a Maxwellian equilibrium can be calculated using the procedure described in Section A.1. The measured damping and oscillation frequencies are plotted in the complex plane along with the zero isocontours of $\mathrm{Im}(\epsilon)$ and $\mathrm{Re}(\epsilon)$, where ϵ is the dispersion function. The intersections of the isocontours indicate solutions to the dispersion relation $\epsilon(\omega, k) = 0$. We observe excellent agreement between theory and our numerical results.

4.6.2 The electrostatic Dory-Guest-Harris instability

As a more challenging benchmark problem, we verify that our code calculates the expected linear growth rate for the electrostatic Dory-Guest-Harris (DGH) instability. The DGH instability was identified and its dispersion relation calculated by Dory et al. [29]. It affects loss-cone distributions; that is, distribution functions which have been evacuated of particles with insufficiently large perpendicular velocity. Such distributions can arise in magnetic mirrors. If the distribution function has a sharp enough peak at large $v_\perp$, it is unstable to transverse $k_\| = 0$ perturbations.

The DGH instability was further developed as a benchmark problem for continuum kinetic Vlasov solvers in [109]. In that work specific wavenumbers are excited along with an azimuthal perturbation that is found to best excite the dominant mode, leading to a very clean "startup". The electromagnetic extension of this work was developed in [23]. Later, the eigenfunctions of the DGH instability were determined [22, 21].

To benchmark our electrostatic Vlasov-Poisson solver, we return to the work reported in [109], and perform a single-species Vlasov-Poisson simulation with two different wavenumbers, one of which excites a purely growing mode, and the other of which excites a mode with complex frequency. The initial condition is selected as follows:

$$f_0(x, v_x, v_z) = \frac{1}{\pi \alpha_\perp^2 j!} \left(\frac{|\boldsymbol{v}_\perp|^2}{\alpha_\perp^2} \right)^j \exp\left(-\frac{|\boldsymbol{v}_\perp|^2}{\alpha_\perp^2} \right) \left(1 + \delta \sin\left(4\theta - \frac{k\omega_c\tau}{v_{\perp 0}} x \right) \right), \tag{4.6.1}$$

where $\theta = \arctan(v_z/v_x)$ is the gyrophase angle and $\alpha_\perp = \sqrt{2/j}$. The integer parameter j determines the "peakedness" of the loss-cone distribution at the most common perpendicular velocity $v_{\perp 0} = \sqrt{2}$. Figure 4.2 reproduces a plot from [109] showing the equilibrium f as a function of $|\boldsymbol{v}_\perp|$ for different values of j. As the plot makes clear, the DGH equilibrium represents a challenging problem for the Hermite spectral representation. The loss cone distribution is quite far from a Maxwellian, and thus requires many Hermite modes to capture.

The perturbed equilibrium (4.6.1) is initialized on a periodic domain of length $L_x = \frac{2\pi v_{\perp 0}}{k\omega_c\tau}$. The fundamental frequencies of the plasma are set to $\omega_p\tau = 1.0$, $\omega_c\tau = 1/20$, and $\nu_p\tau = 0.0$. We use $N_x = 32$ Fourier modes in the x direction, and $N_{vx} = N_{vz} = 96$ Hermite modes in both velocity space dimensions. We consider two wavenumbers in detail: $k = 3.15$, which is a purely growing mode, and $k = 4.65$ which is a complex frequency mode. The simulations are run to $\omega_p t = 500$ and $\omega_p t = 750$ respectively with a timestep of $\Delta t = 0.05$. The electric energy trace from each case, along with a linear fit line, is shown in Figure 4.3. The measured frequencies match the results from [109] to three significant digits in each case.

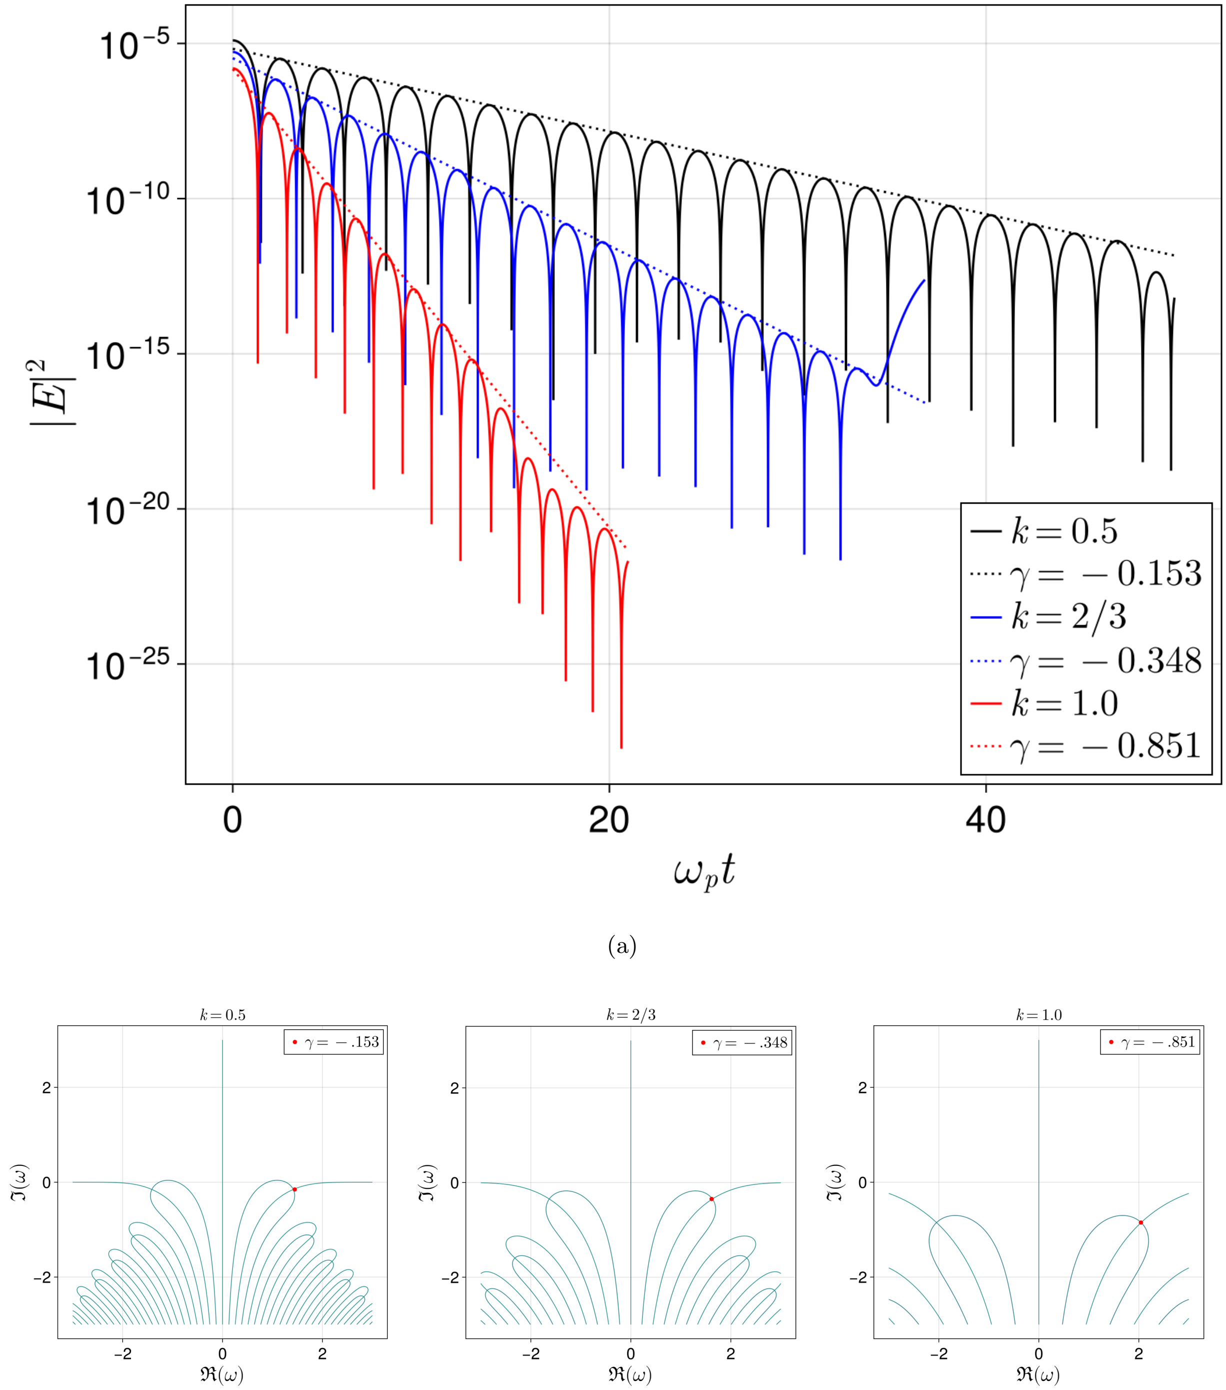

(a)

(b) (c) (d)

Figure 4.1: **Top:** Electric energy traces of three weak perturbations of $\alpha = 1 \times 10^{-3}$, with wavenumbers $k = 0.5, k = 2/3, k = 1.0$. **Bottom:** Zero contours of the real and imaginary parts of the dispersion function $\epsilon(\omega, k)$ for (a) $k = 0.5$, (b) $k = 2/3$, and (c) $k = 1.0$. Red dots mark the location of the complex zero of ϵ with greatest imaginary part γ, which corresponds to the measured damping rate.

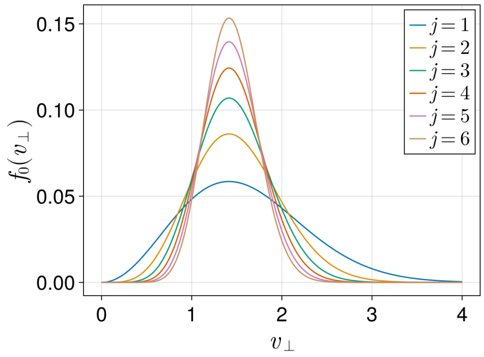

Figure 4.2: The equilibrium phase space density for the Dory-Guest-Harris instability for different values of j. This figure reproduces Figure 1 from [109].

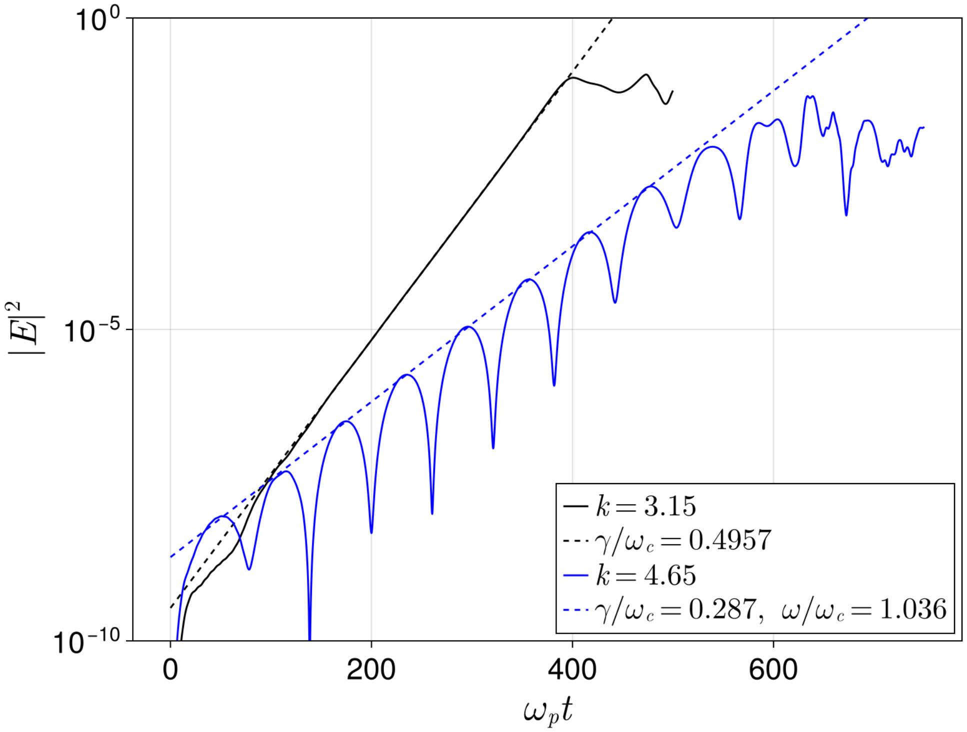

Figure 4.3: Dory-Guest-Harris growth rates for the two wavenumbers considered. The linear phase is measured in the interval $t \in [150, 360]$ for the $k = 3.15$ case and in the interval $t \in [138, 510]$ for the $k = 4.65$ case. This matches the procedure used in [109], Figures 8 and 9.

Chapter 5

Numerical evaluation of finite Larmor radius transport closures

This chapter applies the theory developed in Chapter 3 and the numerical methods described in Chapter 4 to magnetized plasma turbulence. In particular, we consider a family of sheared flow configurations with applied vorticity. Sheared flow is ubiquitous in plasma due to the large variety of drift velocities, including the $E \times B$ velocity, diamagnetic velocity, and various drifts resulting from the magnetic geometry. The resulting sheared flows can give rise to turbulent transport through the mechanism of instability, particularly the Kelvin-Helmholtz instability and subsequent secondary instabilities.

This study is less focused on the specifics of the instability mechanism, and more focused on how particles and heat are transported in the resulting plasma flow, which can exhibit strong velocity shear, compression, and temperature gradients. In the presence of these features, the transport theory of Chapter 3 predicts significant non-ideal transport. To focus on these transport physics, we initialize strong perturbations of a fluid equilibrium, injecting large amounts of vorticity into the flow directly. This is in contrast to typical studies of instability growth, which initialize small perturbations with the intent of studying the linear phase in depth.

Asymptotic theories such as the one developed in Chapter 3 are powerful, but they have severe limitations. Most obviously, they are limited by the maximum order of the expansion to which they are carried out. Thus while we can say with confidence that as $\epsilon \to 0$, eventually the leading-order term will dominate, it is difficult in principle to anticipate how small ϵ will have to be for the theory to be predictive. It is therefore of interest to translate the scaling normalization described in Section 3.1 to a complete initial condition, and observe how the agreement between theory and practice evolves.

This chapter is organized as follows. In Sections 5.1 and 5.2 we describe a family of magnetized initial conditions that produces strong transport phenomena such as heat fluxes and gyroviscosity, while minimizing confounding dynamics such as plasma waves. Section 5.3 reviews numerical results from simulations of these initial conditions and considers the quality of the diamagnetic heat flux and gyroviscous stress tensor approximations from a statistical point of view.

In Section 5.4 we propose and evaluate an adjustment to the Braginskii gyroviscous stress which greatly improves its accuracy in the presence of temperature gradients and retains its ease of numerical implementation. The adjustment accounts for over- or under-prediction of the gyroviscous stress due to the diamagnetic drift velocity. Finally, Section 5.5 examines the residuals of the leading-order approximations and discusses their connection to higher-order terms in the asymptotic expansion. Section 5.6 closes with additional discussion.

5.1 Initial conditions

To more cleanly elicit leading-order kinetic transport physics, we construct an initial condition with an incompressible velocity field. Our aim is to minimize the impact of waves on the solution, since wave phenomena in general will have a shorter timescale than the dynamic timescale on which the vortex flow field evolves. We seek an initial condition satisfying

$$\frac{\mathrm{d}}{\mathrm{d}t}\begin{pmatrix} n_i \\ n_i\boldsymbol{u}_i \\ T_i \end{pmatrix}\Bigg|_{t=0} = 0. \tag{5.1.1}$$

This property gives a “clean” initialization to the simulation, in much the same way as perturbing an equilibrium with a small perturbation provides a clean start. We are concerned only with the ion initial condition here, reasoning that electron inertia is small enough that the electrons will not pollute the solution much at all.

From the non-conservative form of the five-moment equations, we have

$$\frac{\mathrm{d}}{\mathrm{d}t}n_i + n_i\nabla\cdot\boldsymbol{u}_i = 0, \tag{5.1.2}$$

$$A_i\frac{\mathrm{d}}{\mathrm{d}t}(n_i\boldsymbol{u}_i) + \nabla p_i = n_i Z_i(\omega_p\tau\boldsymbol{E} + \omega_c\tau\boldsymbol{u}_i\times\boldsymbol{B}), \tag{5.1.3}$$

$$\frac{\mathrm{d}}{\mathrm{d}t}T_i + (\gamma-1)T_i\nabla\cdot\boldsymbol{u}_i = 0. \tag{5.1.4}$$

From (5.1.2), we have that (5.1.1) will be satisfied as long as

$$\boldsymbol{u}_i = \frac{\omega_p\tau\boldsymbol{E}\times\boldsymbol{B}}{\omega_c\tau|B|^2} + \frac{\boldsymbol{B}\times\nabla p}{\omega_c\tau n_i Z_i|B|^2},$$

and $\nabla\cdot\boldsymbol{u}_i = 0$. The incompressibility condition is automatically satisfied by the $E\times B$ drift. For the diamagnetic drift, we compute

$$\nabla \cdot \left(\frac{\boldsymbol{B} \times \nabla p_i}{n_i} \right) = \nabla \cdot (\boldsymbol{B} \times \nabla T_i + T_i(\boldsymbol{B} \times \nabla(\ln n_i))) \tag{5.1.5}$$

$$= \nabla T_i \cdot (\boldsymbol{B} \times \nabla(\ln n_i)). \tag{5.1.6}$$

Thus, so long as ∇T_i and ∇n_i are colinear, the diamagnetic velocity is incompressible.

We obtain colinearity of the ion density and temperature gradients in the following way. Define an overall ion pressure profile function $\hat{p}_i(z)$, depending only on z, by

$$\hat{p}_i(z) = 1 + \gamma \tanh\left(\frac{z}{\alpha}\right),$$

where α is a length parameter setting the width of the pressure interface. We control the variation of the ion density and temperature together via a parameter $\zeta \in [0, 1]$:

$$n_{i0}(z) = n_{ref}\hat{p}(z)^{\zeta}, \quad T_{i0}(z) = T_{ref}\hat{p}(z)^{1-\zeta}. \tag{5.1.7}$$

The electron species is used to initialize an $E \times B$ drift which has both shear and vorticity. These are controlled by two parameters, u_s and u_V. The first of these represents the desired change in u_x from the bottom to the top of the domain:

$$u_s \triangleq \left.\frac{-\omega_p \tau E_z}{\omega_c \tau B}\right|_{z=L_z/2} - \left.\frac{-\omega_p \tau E_z}{\omega_c \tau B}\right|_{z=-L_z/2}.$$

The latter represents the desired maximum z-directed velocity at the center of the domain:

$$u_V \triangleq \max_x \left.\frac{\omega_p \tau E_x}{\omega_c \tau B}\right|_{z=0}.$$

These $E \times B$ velocities are obtained by prescribing an electrostatic potential $\phi^*(x, z)$, which is given by

$$\phi^*(x, z) = \frac{\omega_c \tau B}{\omega_p \tau} \phi_X^*(x, z)\phi_Z^*(z),$$

where

$$\phi_Z^*(z) = \left(1 + \frac{u_s \alpha}{2} \ln\cosh\left(\frac{z}{\alpha}\right)\right),$$

$$\phi_X^*(x, z) = 1 + \frac{u_V}{k_x} \sin(k_x x) \exp\left(-\frac{z^2}{2w^2}\right).$$

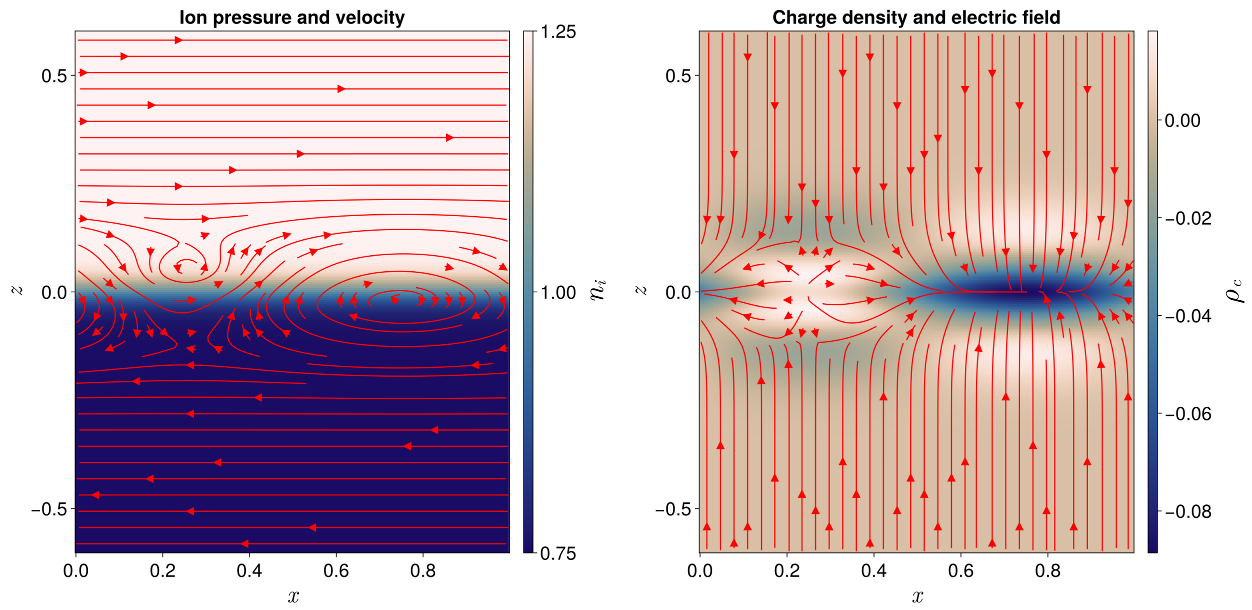

(a) Plot of ion density n_{i0} overlaid with ion velocity $\boldsymbol{u}_{i0}$. (b) Plot of charge density ρ_{c0} overlaid with electric field $\boldsymbol{E}_0$.

Figure 5.1: Illustrative example of an incompressible ion flow initial condition in a domain with sizes $L_x = 1.0$, $L_z = 1.2$. The parameters chosen are $\gamma = 0.25$, $\omega_c\tau = \omega_p\tau = 2.0$, $u_s = 0.2v_{ti}$, $u_V = 0.1v_{ti}$, where $v_{ti} = \sqrt{T_{ref}/A_i}$ is the ion thermal speed. The geometric parameters are $\alpha = 0.04$ and $w = 2\alpha$.

We have introduced two further geometric parameters, k_x and w, which set the wavenumber and width, respectively, of the vorticity.

Given this potential profile, the resulting electric field has components

$$
\begin{aligned}
E_x &= -\partial_x \phi^*(x,z) = -\frac{\omega_c\tau B}{\omega_p\tau}\phi^*_Z(z)\left(u_V\cos(k_x x)\exp\left(-\frac{z^2}{2w^2}\right)\right),\\
E_z &= -\partial_x \phi^*(x,z) = -\frac{\omega_c\tau B}{\omega_p\tau}\left[\frac{u_s}{2}\tanh\left(\frac{z}{\alpha}\right)\phi^*_X(x,z) - \frac{u_V z}{w k_x}\sin(k_x x)\exp\left(-\frac{z^2}{2w^2}\right)\right].
\end{aligned}
$$

Finally, the charge density ρ_c^* which produces this potential can be calculated via $\nabla\cdot\boldsymbol{E} = \omega_p\tau\rho_c^*$:

$$
\begin{aligned}
\rho_c^* = {}& \frac{\omega_c\tau B}{(\omega_p\tau)^2}\phi^*_Z(z)k_x u_V\sin(k_x x)\exp\left(-\frac{z^2}{2w^2}\right)\\
& - \frac{\omega_c\tau B}{(\omega_p\tau)^2}\left[\frac{u_s}{2\alpha}\operatorname{sech}^2\left(\frac{z}{\alpha}\right)\phi^*_X(x,z) - \frac{u_s z}{2w}\sin(k_x x)\exp\left(-\frac{z^2}{2w^2}\right) + \frac{u_V}{k_x}\sin(k_x x)\frac{z^2-w^2}{w^2}\exp\left(-\frac{z^2}{2w^2}\right)\right].
\end{aligned}
$$

The electron density can be determined immediately from the charge density via

$$
n_e = \frac{1}{Z_e}(\rho_c - Z_i n_i).
$$

An example initial condition is plotted in Figure 5.1.

5.2 Kinetic simulations and initial conditions

The fluid parameters of the initial condition described in the previous section determine a Maxwellian kinetic initial condition. However, using a Maxwellian initial condition for the marginally magnetized configurations we explore here is a dubious choice. While we have ensured that the total derivative $\frac{\mathrm{d}}{\mathrm{d}t}$ of each of the Maxwellian fluid parameters vanishes at $t = 0$, this is not the case for the higher-order moments we want to study here. This means that waves of deviatoric pressure and heat flux are excited at early times which make the solution harder to interpret.

To mitigate these effects, we initialize a kinetic ion distribution function whose full pressure tensor and heat flux match the predictions from the asymptotic theory, specifically the hyperbolic expansion summarized in Section 3.5.2. That is, we seek an ion initial condition f_{i0} satisfying

$$\begin{aligned}
\mathbb{P}_{i0} &= A_i \int (\boldsymbol{v} - \boldsymbol{u}) \otimes (\boldsymbol{v} - \boldsymbol{u}) f_{i0}\,\mathrm{d}\boldsymbol{v} = p_{i0}\mathbb{I} + \left[p_{i0}\frac{\mathbb{W}_3[\boldsymbol{u}]}{2\omega_c} + \frac{\mathbb{W}_3[\boldsymbol{q}^1]}{4\omega_c} + \frac{A_i n_{i0}}{2}\widehat{\mathsf{U}}\right], \\
\boldsymbol{q}_{i0} &= \frac{A_i}{2}\int (\boldsymbol{v} - \boldsymbol{u})|\boldsymbol{v} - \boldsymbol{u}|^2 f_{i0}\,\mathrm{d}\boldsymbol{v} = 2p_{i0}\boldsymbol{u}_{Ti},
\end{aligned}$$

where

$$\mathsf{U} = \boldsymbol{u}_{Ti} \otimes (\boldsymbol{u} + \boldsymbol{u}_{Ti})$$

with

$$\boldsymbol{u}_{Ti} = \frac{\hat{\boldsymbol{y}} \times \nabla T_i}{Z_i \omega_c \tau |B|}.$$

The full pressure tensor is easily prescribed with a non-isotropic Maxwellian *ansatz*:

$$\hat{f}_i = \frac{A_i n_{i0}}{2\pi|\mathbb{T}_{i0}|^{1/2}} \exp\left(-\frac{A_i \boldsymbol{w}^T \mathbb{T}_{i0}^{-1} \boldsymbol{w}}{2}\right), \tag{5.2.1}$$

where $\mathbb{T}_{i0} = \frac{1}{n_{i0}}\mathbb{P}_{i0}$ is the temperature tensor, $\boldsymbol{w} = \boldsymbol{v} - \boldsymbol{u}_i$ the relative velocity, $\mathbb{T}_{i0}^{-1}$ is the matrix inverse, and $|\mathbb{T}_{i0}|$ the matrix determinant. That (5.2.1) gives a distribution with the correct pressure tensor can be verified using standard properties of the multivariate Gaussian distribution, viewing $\mathbb{T}_{i0}$ as the covariance matrix.

To additionally prescribe the correct heat flux, we add a component to $\hat{f}_i$ having heat flux $\boldsymbol{q}_{i0}$ and vanishing lower moments. This can be accomplished by defining

$$f_{i0} = \hat{f}_i + \mathcal{M}_{i0}\left[\frac{2\boldsymbol{q}_{i0}}{A_i n_{i0} v_{ti}^3 \sqrt{6}} \cdot \boldsymbol{H}_3\right],$$

where $v_{ti} = \sqrt{T_{i0}/A_i}$,

$$\boldsymbol{H}_3 = \left(He_3\left(\frac{w_x}{v_{ti}}\right), He_3\left(\frac{w_z}{v_{ti}}\right) \right)^T,$$

and

$$\mathcal{M}_{i0} = \frac{A_i n_{i0}}{2\pi T_{i0}} \exp\left(-\frac{A_i|\boldsymbol{w}|^2}{2T_{i0}}\right)$$

is the local Maxwellian. Because the Hermite polynomials in the components of $\boldsymbol{H}_3$ are orthogonal with respect to $\mathcal{M}_{i0}/n_{i0}$, we have

$$\begin{aligned}
\frac{A_i}{2}\int w_x|\boldsymbol{w}|^2 f_{i0}\,\mathrm{d}\boldsymbol{v} &= \frac{q_{i0x}}{n_{i0}}\int \frac{w_x^3 + w_x w_z^2}{v_{ti}^3\sqrt{6}} He_3\left(\frac{w_x}{v_{ti}}\right)\mathcal{M}_{i0}\,\mathrm{d}\boldsymbol{v} \\
&= \frac{q_{i0x}}{n_{i0}}\int He_3\left(\frac{w_x}{v_{ti}}\right)^2 \mathcal{M}_{i0}\,\mathrm{d}\boldsymbol{v} \\
&= q_{i0x},
\end{aligned}$$

and similar for q_{i0z}.

To evaluate the regions of validity of our collisionless, magnetized transport theory, we perform several simulations with varying parameter values. The parameters are summarized in Table 5.1. Series A is designed to explore the role of magnetization in determining the validity of the leading-order transport theory. Magnetization is characterized by the dimensionless parameter $\omega_c\tau$, which in the asymptotic expansion of Chapter 3 is formally connected to the small parameter ϵ. Series A consists of seven simulations with $\omega_c\tau$ varying from 0.5 to 4.5. For reference, the plasma frequency is set by $\omega_p\tau = 1$. Thus, simulation A1 is weakly magnetized relative to electrostatic effects, while simulation A7 is strongly magnetized. Series A fixes the parameter ζ at 0.5, which balances the density and temperature contributions to the pressure gradient (and therefore diamagnetic drift). Thus, the heat flux correction to the gyroviscous stress tensor is expected to be non-negligible in these simulations.

Series B and C are designed to explore the role of temperature gradients in driving the gyroviscous stress. In these series the parameter ζ is varied from -0.5 to 2.0 in increments of 0.5. Per (5.1.7), a value of $\zeta = 0.0$ represents a uniform density profile, while $\zeta = 1.0$ represents an isothermal initial condition. Setting $\zeta = -0.5$ gives a large temperature gradient and a density profile which is oriented opposite the pressure gradient, while $\zeta = 2.0$ gives the reverse: a large density gradient and a temperature gradient oriented opposite the pressure gradient. By varying the relative contribution of temperature to the pressure gradient, we control the relative magnitude of the heat flux and diamagnetic drift in the shear layer, and correspondingly, the relative magnitude of the first two terms of (5.3.1). A larger relative contribution of the second term of (5.3.1) corresponds to larger deviation from the Braginskii gyroviscous stress closure, as discussed in Section

Parameter	Description	A1-A7	B1-B6	C1-C6	M1-M4	S1-S2
$\omega_c \tau$	Magnetization	$\{0.5, 0.75, \ldots, 4.5\}$	2.0	4.0	$\{1.5, 2.0, 3.0, 4.5\}$	2.0
γ	Pressure jump	0.4				0.25
u_s	Shear velocity	$0.2 v_{ti}$				$0.2 v_{ti},\ -0.2 v_{ti}$
u_V	Vortex velocity	$0.06 v_{ti}$			$0.036 v_{ti}$	$0.1 v_{ti}$
ζ	Density/temperature balance	0.5	$\{-0.5, 0.0, \ldots 2.0\}$		0.5	0.5
k_x	Wavenumber	2π			$\{\pi, 2\pi\}$	2π
T_{ref}	Reference temperature	1×10^{-3}				
α	Interface width	0.04				
A_e	Electron mass	1/1836				

Table 5.1: Summary of simulation parameter values.

5.4. In this way we can investigate the importance of temperature gradients in driving gyroviscous transport of momentum. Moreover, series B and C are run with different values of $\omega_c \tau$, with the aim of elucidating the importance of magnetization on the heat flux correction.

To evaluate the role of nonlinear turbulent dynamics in transport closure validity, we run a series of simulations with a superposition of multiple sinusoidal modes in the initial velocity field, series M. We generalize the imposed electrostatic potential by defining

$$\phi_X^*(x, z) = 1 + \sum_i \frac{u_V}{k_x^i} \sin(k_x^i x) \exp\left(-\frac{z^2}{2w^2}\right),$$

for a collection of wavenumbers k_x^i. We apply two modes with wavenumbers $k_x^1 = \pi, k_x^2 = 2\pi$. Additionally we widen the domain to $L_x = 2.0$, and reduce the vorticity velocity u_V compared to series A.

Finally, we seek to understand the role of the polarity of sheared flow in FLR effects. This is accomplished through simulations S1 and S2, which are initialized with opposite shear polarities, defined as the sign of $(\nabla \times \boldsymbol{u}) \cdot \boldsymbol{B}$. The polarity of the sheared flow relative to the magnetic field has been found to impact the linear growth rate of magnetized Kelvin-Helmholtz instabilities [101, 102, 108]. This effect was observed in [108] to be connected to ion inertia through the polarization drift.

Each of the simulations listed in Table 5.1 is performed using an electrostatic hybrid model, which evolves the ion species according to the Vlasov equation, and the electron species according to the drift-advection equation (3.5.1). This hybrid model allows us to avoid using a reduced ion/electron mass ratio by instead stepping over electron kinetic timescales. Unphysically large electron inertia was found in [96] to have a large impact on FLR effects in ten-moment simulations of magnetized Rayleigh-Taylor instabilities. In particular, [96] reports exaggerated ion agyrotropy for reduced mass ratio ten-moment simulations. Because our focus is on accurate evaluation of ion kinetic moments, and because the simplified geometry of our problem allows it, we opt for the simplified drift-advection model for the electron species rather than potentially impacting ion kinetic moments through a reduced mass ratio. Our results are qualitatively similar to ten-moment simulations with a physical mass ratio conducted with the gkylzero code [53], indicating that both the electrostatic and massless electron approximations are accurate for the problems considered here.

All simulations are run on a grid of size $N_x \times N_z = 144 \times 280$ with a CFL number of 0.7. The CFL number is calculated relative to the maximum eigenvalue of the Vlasov operator, which we conservatively estimate as the sum [41] of the eigenvalues corresponding to the free-streaming and Lorentz force operators,

$$\begin{aligned} \lambda_{fs} &= \sum_{\xi=x,y,z} \frac{v_{th}\sqrt{N_{v\xi}}}{\Delta\xi}, \\ \lambda_{Lorentz} &= \sum_{\xi=x,y,z} \frac{\sqrt{N_{v\xi}}}{v_{th}} \left[\omega_p \tau E_\xi + \omega_c \tau v_{th} \sqrt{N_{v\xi}} B_y\right]. \end{aligned}$$

The kinetic solutions use $N_v = 26$ Hermite modes in each dimension, which we find is enough to resolve the ion distribution functions arising in these problems.

5.3 Leading-order convergence

We begin by evaluating the convergence of the closure relations for heat flux and stress tensor as $\omega_c\tau$ increases. The dimensionless magnetization parameter $\omega_c\tau$ is formally connected to the expansion parameter ϵ^{-1}. We therefore expect that agreement between the leading-order closure prediction and observed kinetic moments will improve as magnetization is increased. We recall our derivations from Chapter 3. The gyroviscous stress tensor closure is

$$\Pi \sim \widehat{\Pi}_\perp = p_\perp^0 \frac{\mathbb{W}_3[\boldsymbol{u}_\perp]}{2\Omega_c} + \frac{\mathbb{W}_3[\boldsymbol{q}_\perp]}{4\Omega_c} + \frac{A n_\perp}{2}[(\boldsymbol{u}_\perp^\dagger \widehat{- \boldsymbol{u}_\perp)} \otimes \boldsymbol{u}_\perp] + \frac{A}{2p_\perp^0 T_\perp^0}(\widehat{\boldsymbol{q}_\perp \otimes \boldsymbol{q}_\perp}), \tag{5.3.1}$$

while the heat flux closure is

$$\boldsymbol{q}_\perp \sim \hat{\boldsymbol{q}}_\perp = 2p^0 \frac{\boldsymbol{B} \times \nabla T^0}{Z\omega_c\tau|B|^2}. \tag{5.3.2}$$

The zeroth-order pressure and temperature are defined by $p_\perp^0 = e_\perp,\ T_\perp^0 = p_\perp^0/n_\perp$.

We adopt a statistical perspective with regard to the closures (5.3.1) and (5.3.2), and apply a standard statistical measure of goodness-of-fit, the R^2 value. R^2 is sometimes referred to as the "proportion of variance explained," and it has the following formula for a model $\hat{\psi}$ which is supposed to fit a set of observations ψ:

$$R^2 = 1 - \frac{|\psi - \hat{\psi}|^2}{|\psi - \overline{\psi}|^2}, \tag{5.3.3}$$

where $\overline{\psi}$ is the spatial average of ψ, and the norms are L^2 norms of the spatially dependent functions. For the vector- and tensor-valued transport relations evaluated here, we compute the error and mean componentwise, and then integrate to find the L^2 function norm of the L^2 norm of the error and deviation from the mean. That is, we compute

$$
\begin{aligned}
R^2_\Pi(t) &= 1 - \frac{\left(\int_\Omega |\Pi_\perp(\boldsymbol{x}) - \widehat{\Pi}_\perp(\boldsymbol{x})|^2 \,\mathrm{d}\boldsymbol{x}\right)^{1/2}}{\left(\int_\Omega |\Pi_\perp(\boldsymbol{x}) - \overline{\Pi_\perp}|^2 \,\mathrm{d}\boldsymbol{x}\right)^{1/2}}, \\
R^2_{\boldsymbol{q}}(t) &= 1 - \frac{\left(\int_\Omega |\boldsymbol{q}_\perp(\boldsymbol{x}) - \widehat{\boldsymbol{q}}_\perp(\boldsymbol{x})|^2 \,\mathrm{d}\boldsymbol{x}\right)^{1/2}}{\left(\int_\Omega |\boldsymbol{q}_\perp(\boldsymbol{x}) - \overline{\boldsymbol{q}_\perp}|^2 \,\mathrm{d}\boldsymbol{x}\right)^{1/2}}.
\end{aligned}
$$

These formulas are applied to snapshots of the solution moments. Snapshots are taken of the kinetic simulations by taking a weighted average of f over a time period of length 0.75τ, which for the simulations performed here ranges from approximately one cyclotron period to around 5 cyclotron periods. These weighted averages are then processed by taking moments to obtain $\Pi_\perp$, $\boldsymbol{q}_\perp$, and the inputs to the closures (5.3.1) and (5.3.2). The averaging process smooths over fast variations due to wave phenomena at frequencies close to the cyclotron and plasma frequencies, while leaving the long-time evolution of the moments and transport closures unaffected. In order to better center the snapshots at a point in time, the weighting function is chosen to be a "hat" function which is piecewise linear and symmetric about the point in time to which the snapshot is attributed.

The R^2 values for (5.3.1) are plotted in Figure 5.2 as a function of time for each of the simulations A1-A7 and M1-M4 listed in Table 5.1. They indicate that the accuracy of the transport closure improves significantly as $\omega_c\tau$ increases from 0.5 to 4.5. The similarity of the R^2 traces for $\omega_c\tau = 3.0$ and $\omega_c\tau = 4.5$, however, suggests that further convergence to the leading-order transport theory is beyond the ability of our simulations to discriminate.

Figure 5.3 plots the R^2 values for the heat flux closure (5.3.2). We observe the same overall pattern of improving agreement as $\omega_c\tau$ increases. Notably, the overall trend is that the R^2 for heat flux is higher than the R^2 for the stress tensor, despite the stress tensor closure being formally of order ϵ^2. The results for single-mode (series A) and double-mode (series M) vorticity are quite comparable. In general agreement is better for the two-mode series M simulations, which have an x scale $L_x = 2.0$ of twice that of series A, and thus longer gradient scale lengths in general. At late times such as 500τ and above, all simulations have become highly distorted and begun the transition to turbulent mixing. Figure 5.4 plots the density and temperature of cases M1 and M4 at $t = 600\tau$. The vortex structure is significantly more coherent at this late time for the $\omega_c\tau = 4.5$ case.

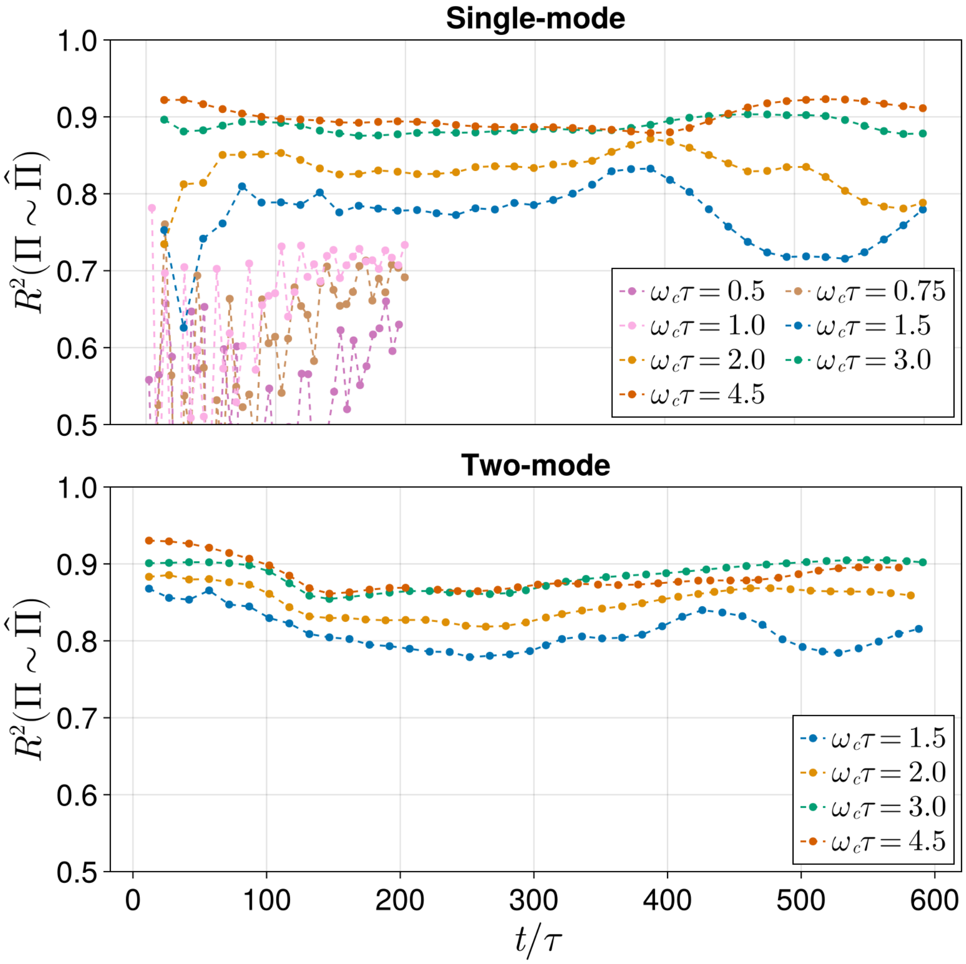

Figure 5.2: R^2 values for (5.3.1) as a function of time for simulations A1-A7 (top) and M1-M4 (bottom).

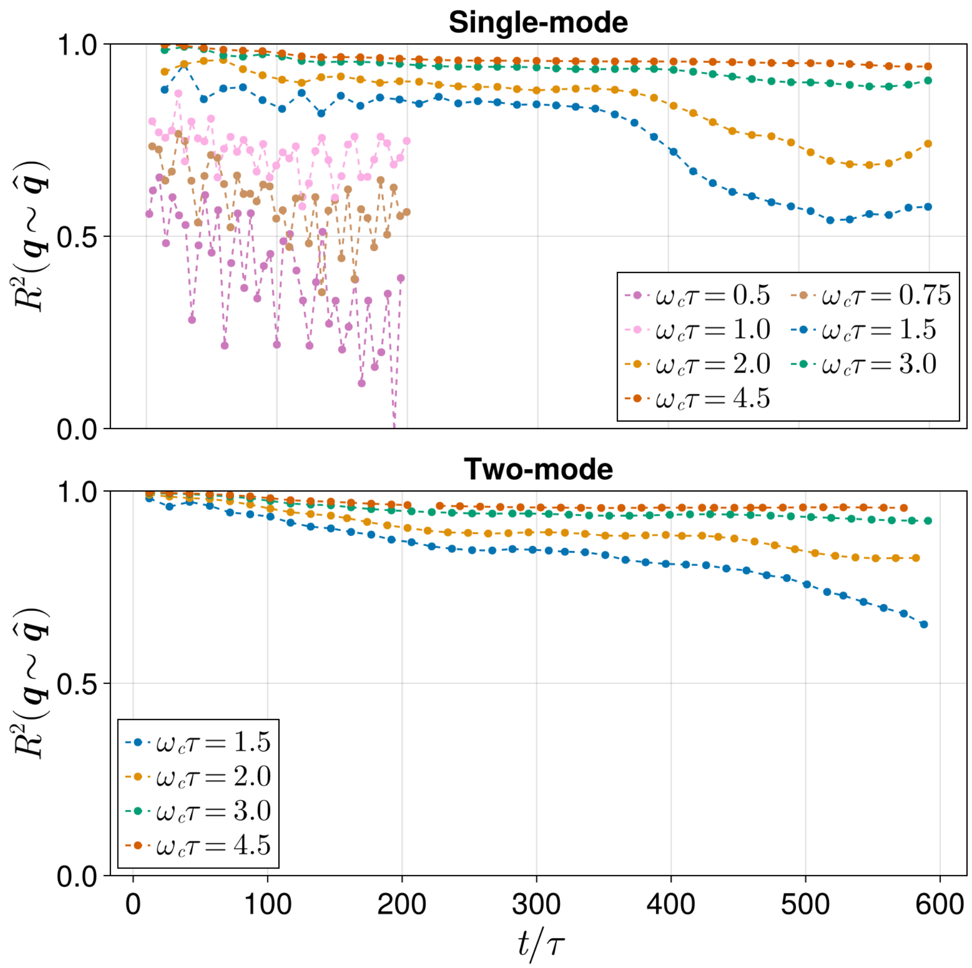

Figure 5.3: R^2 values for (5.3.2) as a function of time for simulations A1-A7 (top) and M1-M4 (bottom).

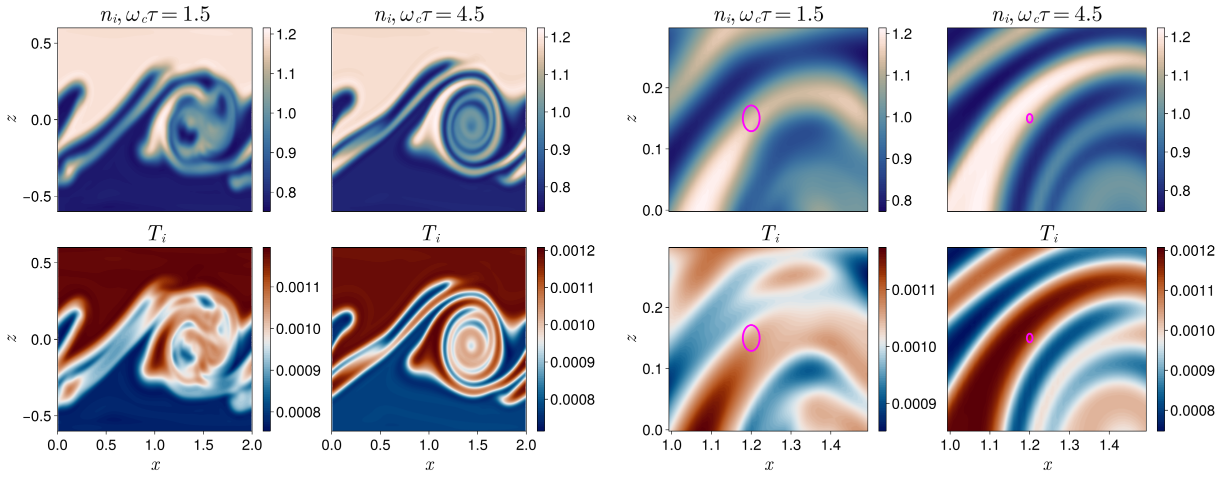

Figure 5.4: **Left:** Density and temperature contours for cases M1 (first column) and M4 (second column), showing late-time vortex structures and the transition to turbulent flow. **Right:** Zoom of the region $(x, z) \in [1.0, 1.5] \times [0.0, 0.3]$ with ellipses indicating characteristic ion Larmor radius $r_{Li} = A_i\sqrt{T_{ref}}/(Z_i\omega_c)$.

5.4 Numerically affordable approximation to leading-order gyroviscous stress

As we have seen, the drift ordering gyroviscous stress tensor closure (5.3.1) is predictive for a range of magnetizations. However, it is difficult to make use of in practice. To begin with, (5.3.1) depends on the heat flux, which is another unclosed moment variable in the five-moment setting. Substituting the diamagnetic heat flux closure for $\boldsymbol{q}$ in (5.3.1) gives a usable closure, but one that depends on second-order derivatives of temperature. When used in the momentum flux term, the result is a third-order PDE. Such equations are challenging to solve numerically, since they impose more restrictive timestep limits for explicit solvers and may introduce challenges related to dispersive discretization errors.

These considerations motivate us to seek an adjustment to the Braginskii gyroviscous stress which captures the main effects of temperature gradients, but without changing the nature of the PDE to be solved numerically. The Braginskii stress tensor closure is

$$\hat{\Pi}_{\perp}^{Brag} = p_{\perp} \frac{\mathbb{W}_3[\boldsymbol{u}_{\perp}]}{2\omega_c}. \tag{5.4.1}$$

Equation (5.4.1) is derived in [11] for ion transport in a simple plasma. It is also given by Roberts and Taylor in [89] for finite Larmor radius MHD. The Braginskii gyroviscous stress is commonly used in theoretical and computational modeling of ion FLR effects [63, 64, 107, 81].

Compared to (5.4.1), the drift-ordered gyroviscous stress closure (5.3.1) predicts modified gyroviscous stress depending on the orientation of the temperature gradient. The mechanism is non-Maxwellian distortion of the distribution function by the diamagnetic heat flux in the presence of temperature gradients. In this way, the diamagnetic heat flux impacts particle transport, despite not causing diffusive heat transport on its own (being orthogonal to ∇T.)

To estimate the impact of the diamagnetic heat flux on gyroviscous stress, assume that the velocity is composed solely of the diamagnetic drift velocity. Substituting the diamagnetic heat flux closure into (5.3.1) gives the estimate

$$\widehat{\Pi}_{\perp} \approx \frac{p_{\perp}^0}{2\omega_c \tau} \mathbb{W}_3 \left[\frac{\boldsymbol{B} \times \nabla p}{n Z \omega_c \tau |B|^2} + \frac{\boldsymbol{B} \times \nabla T}{Z \omega_c \tau |B|^2} \right] + \mathcal{O}(An|\boldsymbol{u}|^2), \tag{5.4.2}$$

where we have collected terms of the magnitude of the kinetic energy. The form of the first two bracketed terms shows that we should predict enhanced (resp. diminished) gyroviscous stress when the temperature gradient is oriented in the same direction (resp. opposite) of the pressure gradient. We proceed to neglect the terms scaling with kinetic energy, which is justified in this regime—flow velocities are of the order of $0.1 v_{ti}$, so kinetic energy is roughly 1% of the internal energy, whereas the stress tensor achieves a norm of around 10% the scalar pressure. Therefore, the following estimates apply:

$$|\widehat{\Pi}_{\perp}^{Brag}| \sim p_{ref} r_{Li} \left|\frac{\nabla p}{p}\right|, \quad \widehat{\Pi}_{\perp} \sim p_{ref} r_{Li} \left|\frac{\nabla p}{p} + \frac{\nabla T}{T}\right|.$$

In this simple case, we expect the Braginskii gyroviscosity to be off by a factor of

$$\gamma^{Brag} \triangleq \frac{|\Pi_{\perp}|}{|\widehat{\Pi}_{\perp}^{Brag}|} \approx \frac{|\nabla p/p + \nabla T/T|}{|\nabla p/p|}. \tag{5.4.3}$$

We perform two series of simulations to evaluate the importance of the heat flux-related terms in (5.3.1) and the performance of the Braginskii gyroviscous stress tensor closure when the relative magnitude of the temperature gradient to pressure gradient is modified. Series B fixes $\omega_c\tau = 2.0$ and varies ζ, which controls the relative magnitude of density and temperature gradients, from $\zeta = -0.5$ to $\zeta = 2.0$. These correspond to values of γ^{Brag} at the initial condition ranging from 2.5 to 0.0. Series C takes the same range of values of ζ, but doubles $\omega_c\tau$ to 4.0.

Figure 5.5 plots $\hat{\Pi}_{xx}^{Brag}$ versus Π_{xx} for simulations C1-C6 at time $t = 200\tau$. The results indicate that the Braginskii closure indeed exhibits the predicted error (5.4.3); in fact, for the plotted simulations the simple adjustment $\Pi \sim \gamma^{Brag}\hat{\Pi}^{Brag}$ is an excellent estimate as demonstrated by the slopes of the black lines of best fit.

These considerations lead us to propose the following modification to the Braginskii ion gyroviscous stress:

$$\widehat{\Pi}_{\perp}^{Adj} = \frac{p}{2\omega_c}\left[\gamma^{Brag}\mathbb{W}_3[\boldsymbol{u} - \boldsymbol{u}_E] + \mathbb{W}_3[\boldsymbol{u}_E]\right] \tag{5.4.4}$$

where $\boldsymbol{u}_E$ is the $E \times B$ velocity

$$\boldsymbol{u}_E = \frac{\omega_p\tau \boldsymbol{E} \times \boldsymbol{B}}{\omega_c\tau|B|^2}.$$

This *ansatz* is based on the assumption that the ion velocity is dominantly composed of the $E \times B$ and diamagnetic drifts, and that ion inertial effects, such as those that contribute to the polarization drift, are relatively negligible. Whereas the $E \times B$ drift velocity is tractable to determine in a numerical computation, terms such as the diamagnetic and polarization drifts are much more complicated, because they require estimating a gradient or time derivative of a primary variable. The parameter γ^{Brag} is defined as in (5.4.3). It may be chosen as a global parameter for the whole simulation, or it may be determined based on local gradient scale lengths of pressure and temperature.

Figure 5.6 plots the same data as Figure 5.5, but the x axis has been changed to the adjusted Braginskii gyroviscous stress closure (5.4.4). Comparing the two plots, we observe that the adjustment greatly improves predictive accuracy of the closure, particularly for $\gamma^{Brag} > 1.0$. In these three cases, the adjustment fixes

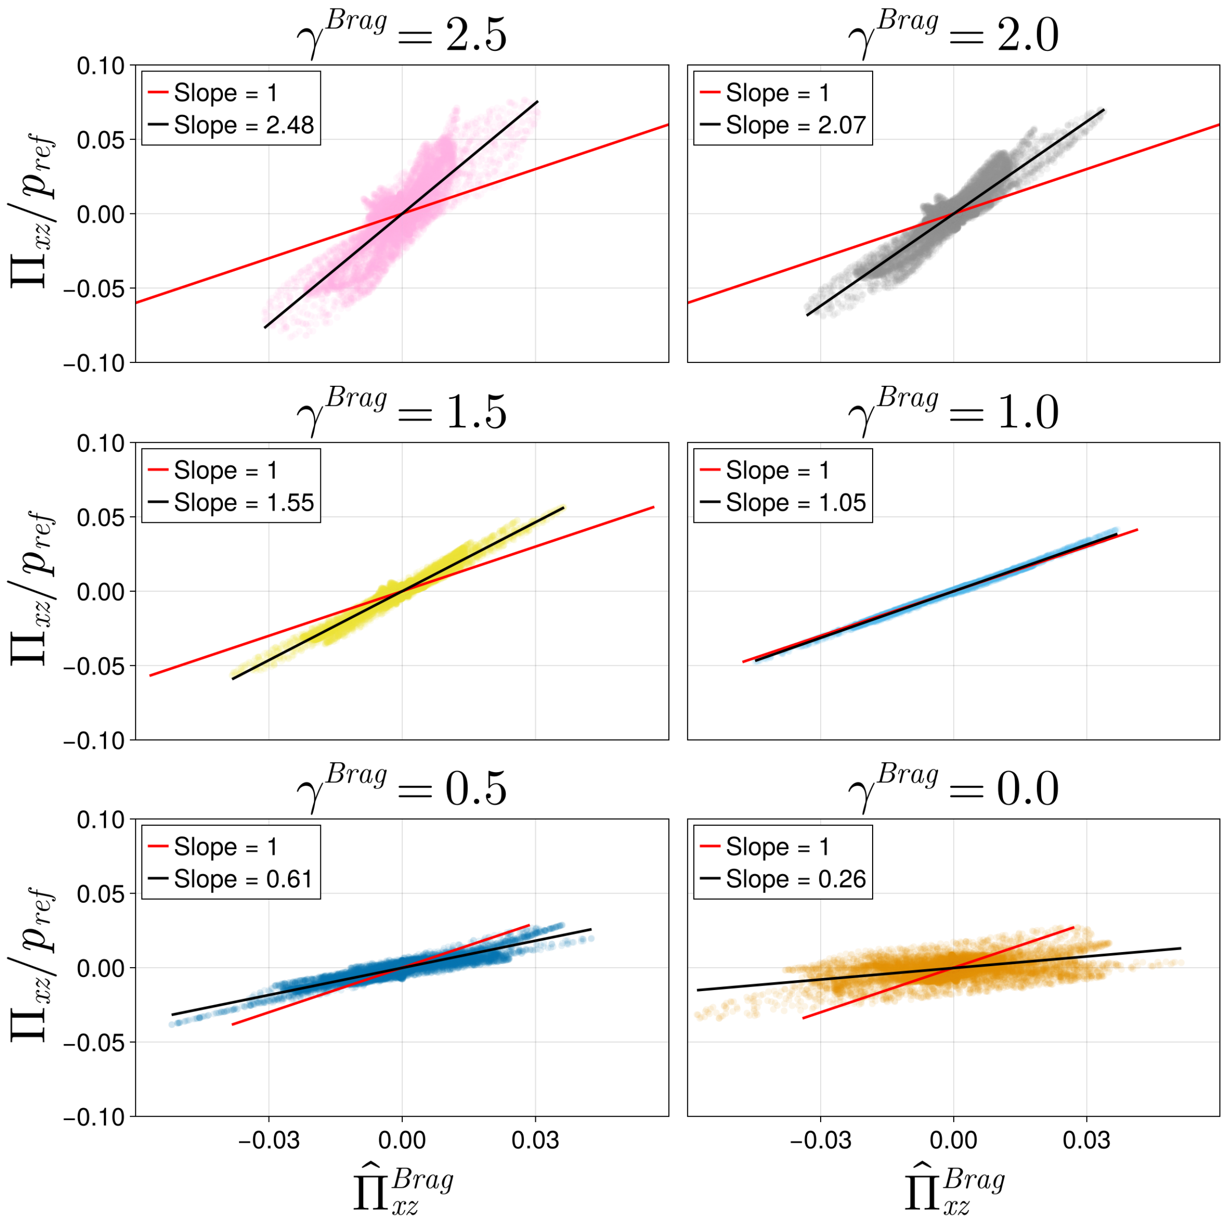

Figure 5.5: Plots of the Braginskii gyroviscous stress tensor prediction (x-axis) versus the observed kinetic stress tensor (y-axis). Black lines are lines of best fit, which have the indicated slopes. Best fit slopes show that $\Pi \approx \gamma^{Brag}\hat{\Pi}^{Brag}$ is a decent approximation across a range of values of γ^{Brag}. Disagreement is most dramatic in the case $\gamma^{Brag} = 0$, where the bottom right scatterplot shows that the Braginskii closure greatly overestimates the magnitude of gyroviscous stress. Simulation data are taken from cases C1-C6, which use $\omega_c\tau = 4.0$, at $t = 200\tau$. Plotted values are from the Π_{xz} component; other components show the same pattern.

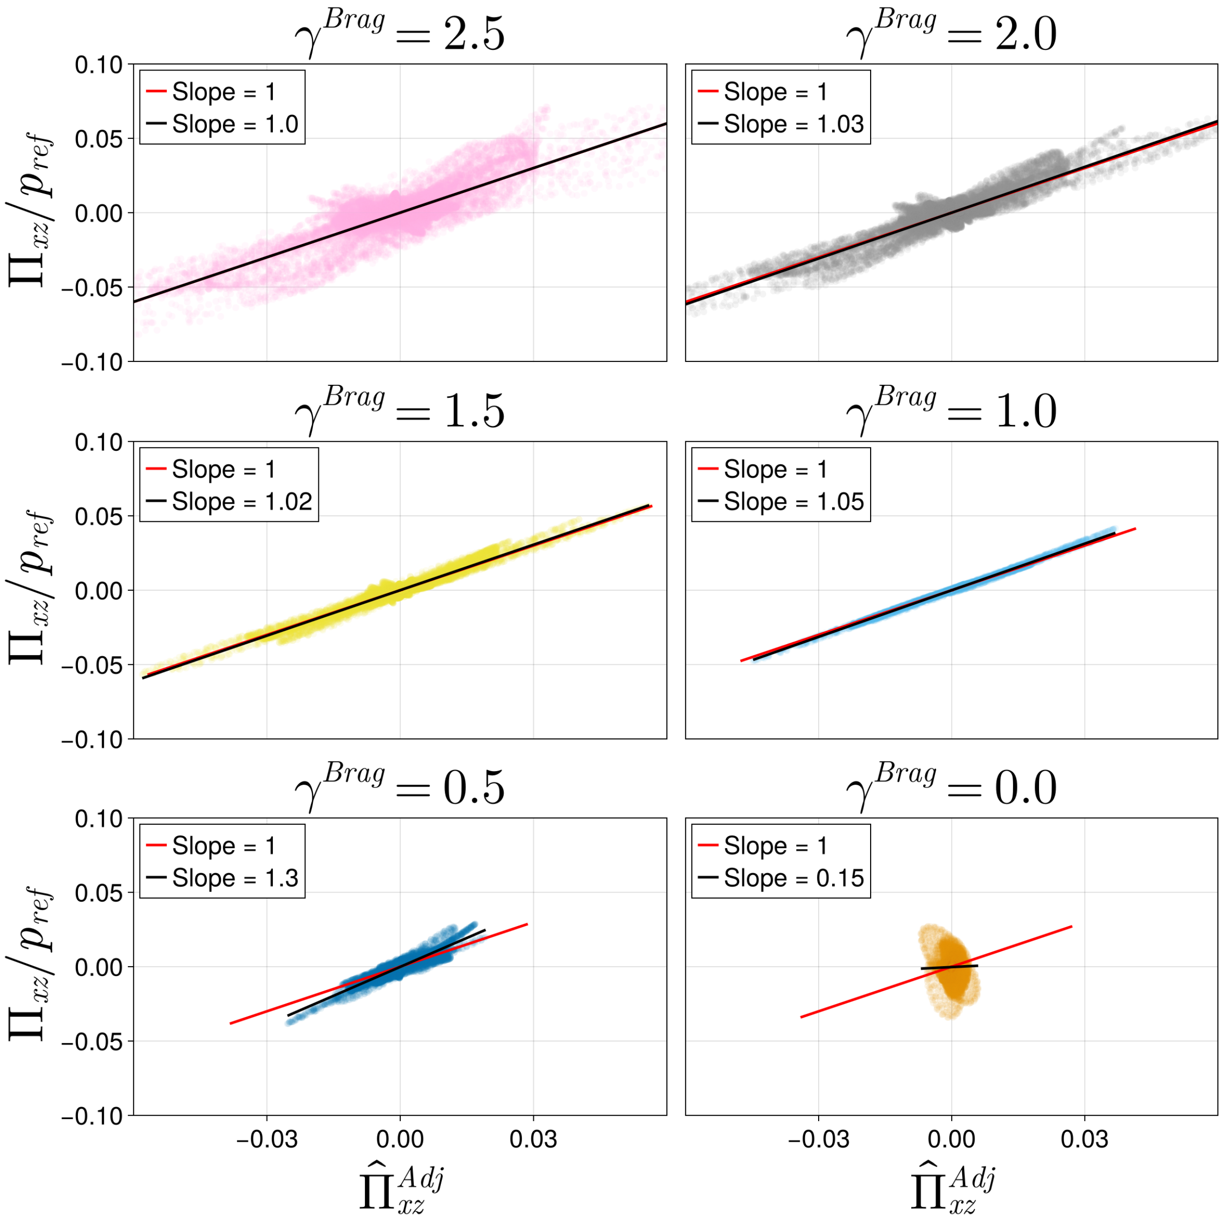

Figure 5.6: Plots of the adjusted Braginskii gyroviscous stress tensor prediction (5.4.4), (x-axis), versus the observed kinetic stress tensor (y-axis). Black lines are lines of best fit, which have the indicated slopes. The simulation data are taken from cases C1-C6, which use $\omega_c\tau = 4.0$, at $t = 200\tau$. Plotted values are from the Π_{xz} component; other components show the same pattern.

systematic underprediction of the gyroviscous stress magnitude, with the line of best fit having a slope much closer to 1.0 than given by the classical Braginskii closure. For the two cases with $\gamma^{Brag} < 1.0$, the adjustment fixes the systematic overprediction of the magnitude of the gyroviscous stress.

Figure 5.7 plots the R^2 value for (5.4.4) over time for cases B1-B4 and C1-C4, all of which have an initial temperature gradient aligned with the pressure gradient. The results indicate that the adjusted Braginskii gyroviscous stress is a greatly improved predictor of Π compared to the unmodified Braginskii gyroviscous stress in the presence of temperature gradients. The quality of the adjusted Braginskii closure is greater for more mild temperature gradients, in particular the case of $\zeta = 0.5, \gamma^{Brag} = 1.5$, which is when density and temperature gradient scale lengths are equal. We note that in this formulation, a uniform density profile is indicated by $\zeta = 0.0, \gamma^{Brag} = 2.0$, for which we also observe significant improvement of the closure. The quality of the adjusted closure begins to drop for $\gamma^{Brag} = 2.5$, which corresponds to a density gradient anti-aligned with the temperature gradient.

The cases with anti-aligned temperature gradients, namely $\gamma^{Brag} = 0.5$ and $\gamma^{Brag} = 0.0$, do not exhibit an improved R^2 value compared to the classical Braginskii closure. This can be observed by comparing the final row of Figures 5.5 and 5.6. In particular, the case of $\gamma^{Brag} = 0.0$ has a much smaller gyroviscous stress magnitude than predicted by Braginskii. The adjustment fixes the overprediction, but the remaining gyroviscous stress appears to be due to terms at higher than leading order, which are not accounted for by either closure.

The adjusted Braginskii closure is an improvement on (5.4.1) across a range of magnetizations and for long-time flow evolution. Figure 5.8 plots the R^2 value for (5.4.1) and (5.4.4) over time for cases A1-A7 and M1-M4. All these cases use $\zeta = 0.5, \gamma^{Brag} = 1.5$. The results show a significant jump in explained variance for $\omega_c\tau \in [1.5, 4.5]$, and less dramatic improvement for $\omega_c\tau \in \{0.75, 1.0\}$. Case A1 with $\omega_c\tau = 0.5$ shows no improvement from the adjusted closure.

5.5 Higher-order corrections in ϵ

The transport closures we have examined so far have been only leading-order closures in the small parameter ϵ, which can be characterized as the ratio of the ion Larmor radius to gradient scale lengths. A complete account of kinetic effects, however, naturally requires terms of order ϵ^2 and higher. We expect that such terms are significant when ϵ is insufficiently small, resulting in deviation of the kinetic heat flux and stress tensor from the leading-order closures. This deviation manifests as reduced R^2 for the corresponding closure, as can be seen in the late-time portion of Figures 5.2, 5.3, and 5.7, as well as in the relatively poor agreement of the closure models for low magnetization.

In the collisionless, magnetized limit, higher-order corrections to the heat flux and stress tensor are of particular interest because the leading-order closure moments do not contribute to diffusion of heat and

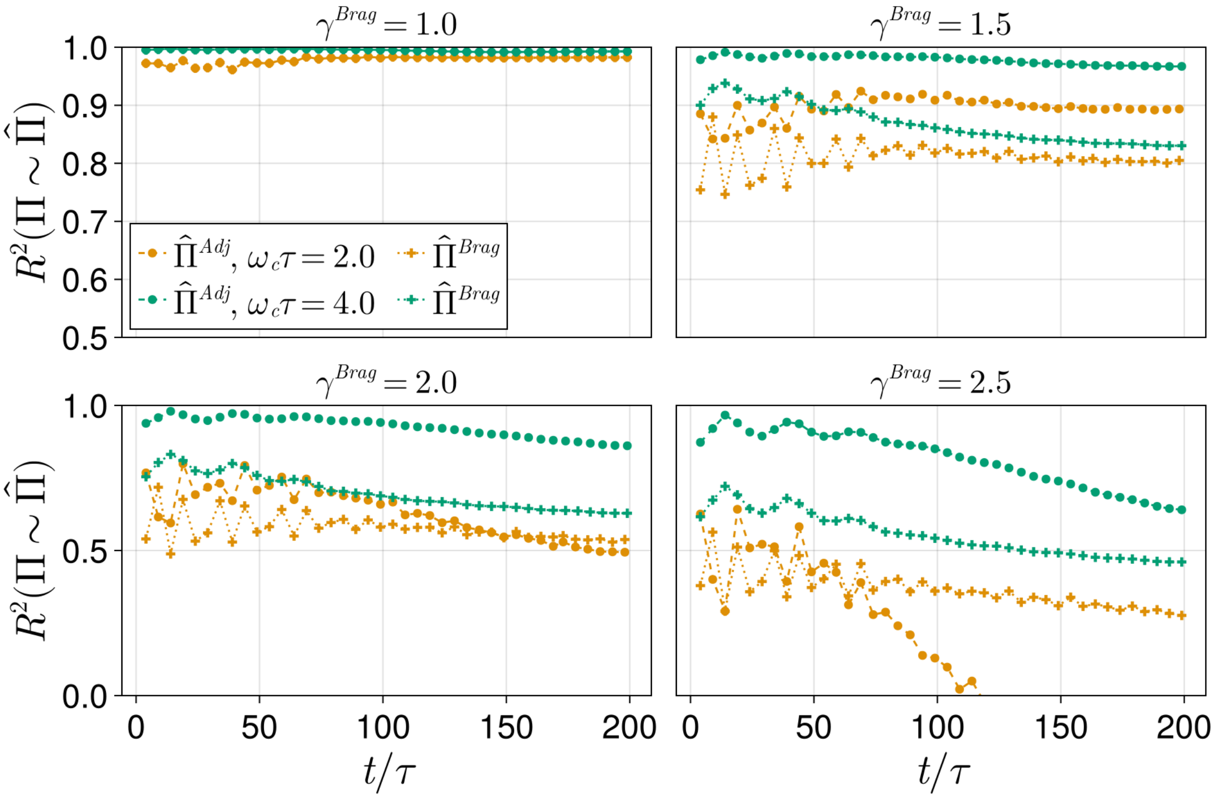

Figure 5.7: Comparison of R^2 values for unmodified (5.4.1) (dotted lines) and adjusted (5.4.4) (dashed lines) gyroviscous stress tensor closures. Plotted simulations are cases B1-B4 ($\omega_c\tau = 2.0$) and C1-C4 ($\omega_c\tau = 4.0$). The late-time results show dramatic improvement of the adjusted closure $\hat{\Pi}_\perp^{Adj}$ over the unmodified Braginskii closure for the high-magnetization case $\omega_c\tau = 4.0$. Agreement of $\hat{\Pi}^{Adj}$ is best for $\gamma^{Brag} = 1.5$ which corresponds to $\zeta = 0.5$, i.e. equal density and temperature gradient scale lengths.

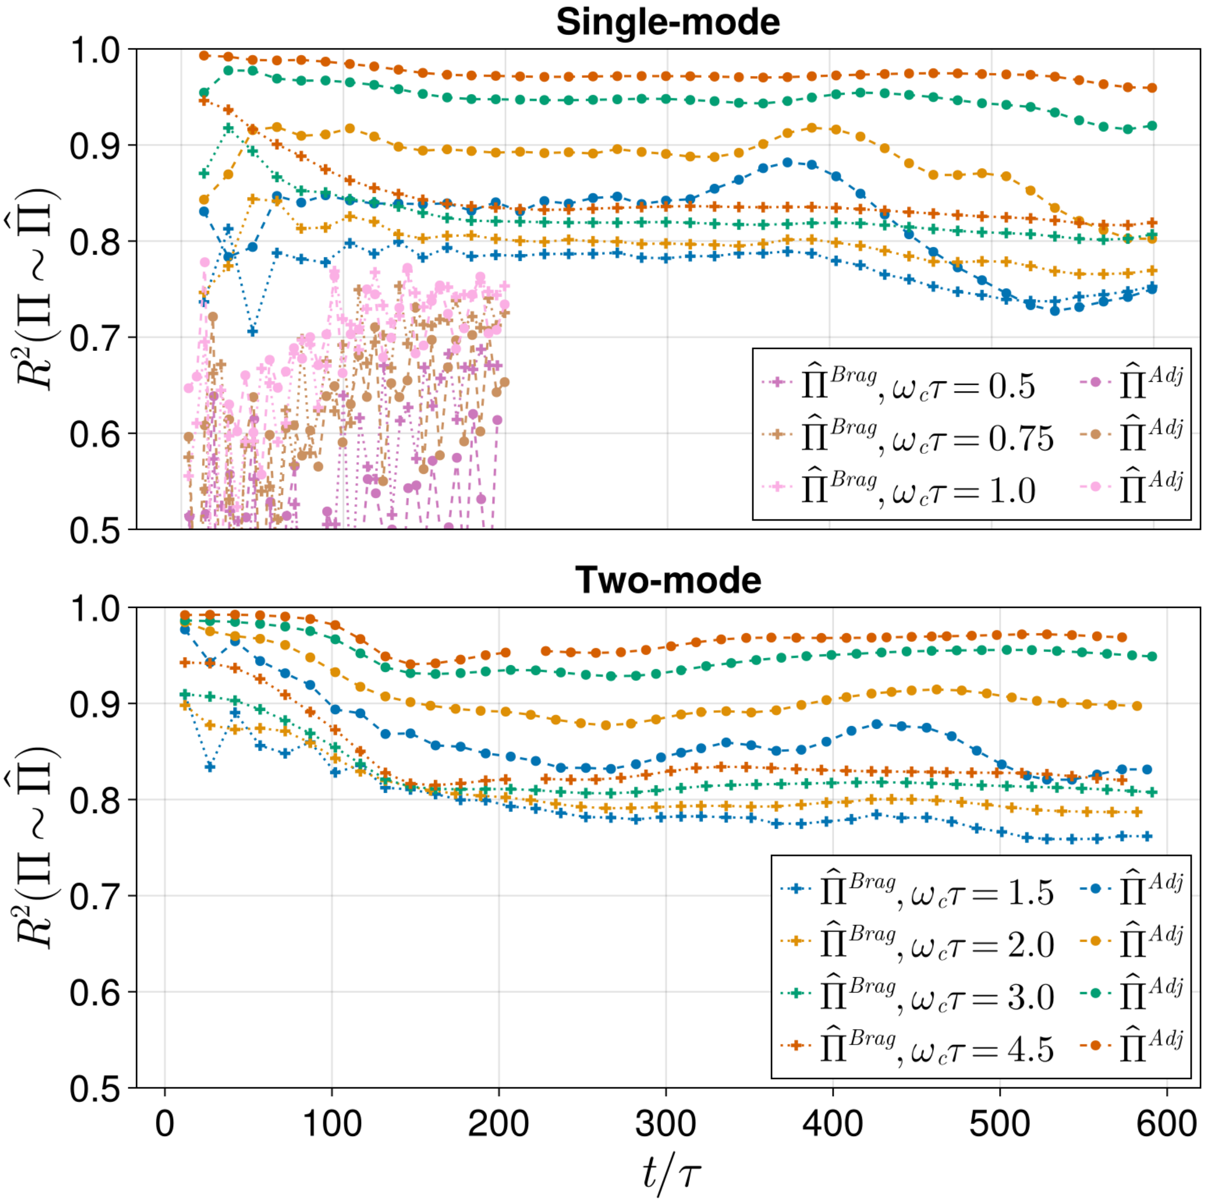

Figure 5.8: Comparison of R^2 values for unmodified Braginskii (5.4.1) (dotted lines) and adjusted (5.4.4) (dashed lines) gyroviscous stress tensor closures. Plotted simulations are cases A1-A7 (top) and M1-M4 (bottom), all of which have $\zeta = 0.5, \gamma^{Brag} = 1.5$.

dissipative viscous heating, respectively. For the heat flux, it is simple to see that

$$\hat{\boldsymbol{q}} \cdot \nabla T = p \frac{\boldsymbol{B} \times \nabla T}{2\Omega_c} \cdot \nabla T = 0,$$

so the diamagnetic heat flux does not transport heat along the temperature gradient. The gyroviscous stress has a similar property, which is revealed by the non-conservative form of the temperature equation:

$$\frac{\mathrm{d}}{\mathrm{d}t} T + \frac{\gamma - 1}{n} (\mathbb{P} : \nabla \boldsymbol{u} + \nabla \cdot \boldsymbol{q}) = 0. \tag{5.5.1}$$

A simple calculation shows that

$$\mathbb{W}_3[\boldsymbol{u}] : \nabla \boldsymbol{u} = 0,$$

which means that the Braginskii gyroviscous stress tensor $\hat{\Pi}_{\perp}^{Brag}$ does not contribute to dissipative viscous heating.

To better understand the role of higher-order corrections to the closure moments, we plot the residuals of the leading-order closures. The heat flux at each point $\boldsymbol{x}$ can be decomposed into a diamagnetic component orthogonal to ∇T and a perpendicular component which is parallel to ∇T. We define the normalized component decomposition of the residual $\boldsymbol{q}_i - \hat{\boldsymbol{q}}_i$ in the following way:

$$[\boldsymbol{q} - \hat{\boldsymbol{q}}]_{\wedge} = \frac{(\boldsymbol{q}_i - \hat{\boldsymbol{q}}_i) \cdot (\boldsymbol{B} \times \nabla T_i)}{A_i |B| |\nabla T_i| v_{ti}^{3/2}}, \quad [\boldsymbol{q} - \hat{\boldsymbol{q}}]_{\perp} = \frac{(\boldsymbol{q}_i - \hat{\boldsymbol{q}}_i) \cdot \nabla T_i}{A_i |\nabla T_i| v_{ti}^{3/2}}. \tag{5.5.2}$$

These expressions are normalized by the free-streaming heat flux limit, which is $A_i v_{ti}^{3/2}$. Figure 5.9 plots these expressions along with the ion temperature for case M3 at four different times. The residual plots reveal coherent structure well into the nonlinear phase. Comparing the residuals with plots of temperature indicate that the heat flux closure residual aligns with regions of high curvature of T_i, consistent with the residual being well-described by a second or third-degree derivative polynomial in T_i. Such expressions arise at higher order in the asymptotic expansion procedure. However, the size of such higher-order corrections naturally has a quadratic or cubic dependence on the inverse temperature gradient scale length. Turbulent flow, being characterized by high-wavenumber spatial features in density and temperature, therefore cannot be expected to conform to the leading-order closure expressions. Moreover, the presence of spatially localized features in the residual plots highlights the importance of using local estimates for closure applicability, rather than global correction factors based on a single problem parameter.

Examining the third row of Figure 5.9 in more detail, we note that the sign of the perpendicular (along-gradient) heat flux residual has a clear dependence on the slope of the rollup in the $x - z$ plane. Case S1, which has a reversed shear direction, is plotted in Figure 5.10 and shows the same trend. A negative sign

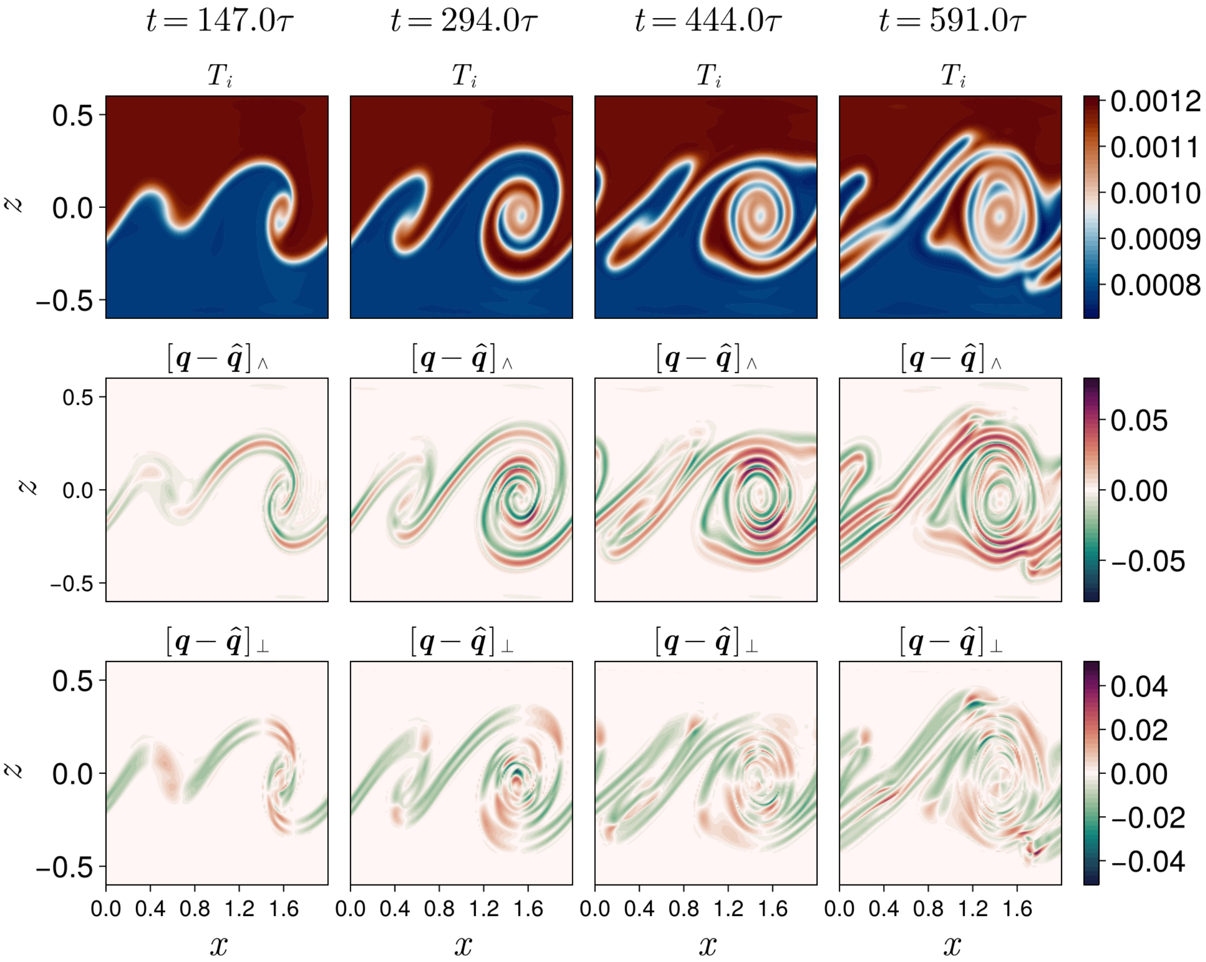

Figure 5.9: Contour plots of the normalized components of the residual $\boldsymbol{q}_i - \hat{\boldsymbol{q}}_i$ for case M3. **First row**: contours of ion temperature. **Second row**: contours of $[\boldsymbol{q} - \hat{\boldsymbol{q}}]_\wedge$, defined as the component of $\boldsymbol{q}_i - \hat{\boldsymbol{q}}_i$ in the direction of $\boldsymbol{B} \times \nabla T_i$. Contours plotted in units of the reference free-streaming heat flux $\boldsymbol{q}_{fs} = v_{ti} p_{ref}$. Spatial structure in the shear layer, in particular multiple sign changes crossing the shear layer, indicate that second- or third-order derivative polynomials of temperature likely play a role in the leading-order residual. This is consistent with the structure of the asymptotic expansions at higher order in ϵ. **Third row**: contours of $[\boldsymbol{q} - \hat{\boldsymbol{q}}]_\perp$, defined as the component of $\boldsymbol{q}_i$ in the direction of ∇T_i. Contours plotted in units of $\boldsymbol{q}_{fs} = v_{ti} p_{ref}$ Negative values of this component indicate diffusion of heat while positive values indicate anti-diffusion of heat. The sign of diffusion exhibits a clear dependence on the slope of the vortex in the $x - z$ plane, consistent with ion inertial effects (see discussion).

of $[\boldsymbol{q}-\hat{\boldsymbol{q}}]_\perp$ indicates a heat flux vector in the opposite direction of ∇T, and thus perpendicular diffusion of heat. On the other hand, a positive sign of $[\boldsymbol{q}-\hat{\boldsymbol{q}}]_\perp$ indicates anti-diffusion of heat.

We hypothesize that this is attributable to higher-order corrections to heat flux associated with ion inertia: as the slope of the vortex increases, the heat flux vector, which is initially in the negative x direction, lags behind the changing $\boldsymbol{B}\times\nabla T$ direction and acquires nonzero components directed parallel to ∇T. To see the origin of this effect, we write the perpendicular moment equation for the heat flux in non-conservative form with index notation:

$$\frac{\mathrm{d}}{\mathrm{d}t}q_i+\partial_j H_{ij}+\partial_j u_j q_i=\frac{Z}{A}\left[\left(\omega_p\tau E_j+\omega_c\tau\epsilon_{jmn}u_m B_n\right)\left[\delta_{ij}p_\perp+\mathbb{P}_{ij}\right]+\omega_c\tau\epsilon_{imn}q_m B_n+\omega_c\tau\epsilon_{jmn}B_n Q_{ijm}\right], \tag{5.5.3}$$

where

$$H_{ij}=\frac{A}{2}\left\langle w_i w_j w_k w_k f\right\rangle_v,\quad Q_{ijm}=\left\langle w_i w_j w_m f\right\rangle_v.$$

Here we are using the notation $\boldsymbol{w}=\boldsymbol{v}_\perp-\boldsymbol{u}_\perp$ and $\langle\cdot\rangle_v=\int\cdot\,\mathrm{d}\boldsymbol{v}_\perp$. Substituting the Maxwellian moments $H_{ij}=\frac{2pT}{A}\delta_{ij}$, $Q_{ijm}=0$, and $\mathbb{P}_{ij}=p_\perp\delta_{ij}$, we can simplify the expression and rewrite in vector notation:

$$\frac{\mathrm{d}}{\mathrm{d}t}\boldsymbol{q}_\perp+\frac{2}{A}\nabla(p_\perp T_\perp)+(\nabla\cdot\boldsymbol{u}_\perp)\boldsymbol{q}_\perp=\frac{Z}{A}\left[(\omega_p\tau\boldsymbol{E}+\omega_c\tau\boldsymbol{u}_\perp\times\boldsymbol{B})\cdot[2p_\perp\mathbb{I}]+\omega_c\tau\boldsymbol{q}_\perp\times\boldsymbol{B}\right]. \tag{5.5.4}$$

Substituting the leading-order drift velocity $\boldsymbol{u}_\perp^1$ into (5.5.4) and neglecting flow compressibility, we get an equation for the leading-order heat flux,

$$2\nabla(p_\perp T_\perp)=Z\left(2\frac{\nabla p_\perp}{n_\perp}p_\perp+\omega_c\tau\boldsymbol{q}_\perp^1\times\boldsymbol{B}\right),$$

whose solution is the diamagnetic heat flux (5.3.2) up to order ϵ^2. At the subsequent order, we substitute the polarization drift $\boldsymbol{u}_\perp^p$ on the right-hand side of (5.5.4) and obtain

$$\frac{\mathrm{d}}{\mathrm{d}t}\boldsymbol{q}_\perp^1=\frac{Z}{A}\left[(\omega_c\tau\boldsymbol{u}_\perp^p\times\boldsymbol{B})(2p_\perp)+\omega_c\tau\boldsymbol{q}_\perp^p\times\boldsymbol{B}\right],$$

or

$$\boldsymbol{q}_\perp^p=\frac{A}{Z\omega_c\tau|B|^2}\left(\frac{\mathrm{d}}{\mathrm{d}t}\boldsymbol{q}_\perp^1\right)\times\boldsymbol{B}-2p_\perp\boldsymbol{u}_\perp^2,$$

The next-to-leading order ion heat flux, which we denote $\boldsymbol{q}_\perp^p$ by analogy with the polarization drift, is therefore seen to be associated with ion inertial effects via the time derivative of the leading-order heat flux as well as the ion polarization drift. Polarization drifts were observed to drive charge accumulation in the Kelvin-Helmholtz instability [108]. The simulations conducted here suggest that similar physics may drive heat accumulation, via along-gradient heat fluxes, in Kelvin-Helmholtz-like vortex structures.

For the stress tensor closure, we split the residual $\Pi_\perp - \hat{\Pi}_\perp$ into two components, one in the direction of $\mathbb{W}_1[\boldsymbol{u}]$ and the other in the direction of $\mathbb{W}_3[\boldsymbol{u}]$. Note that $\mathbb{W}_1[\boldsymbol{u}]$ is defined as

$$\mathbb{W}_1[\boldsymbol{u}] = \begin{pmatrix} \partial_x u_x - \partial_z u_z & \partial_x u_z + \partial_z u_x \\ \partial_z u_x + \partial_x u_z & \partial_z u_z - \partial_x u_x \end{pmatrix},$$

and satisfies $\mathbb{W}_1 : \mathbb{W}_3 = 0$. Since Π is a symmetric, trace-free tensor, it has two degrees of freedom and is therefore uniquely determined by its magnitude in the direction of $\mathbb{W}_1$ and $\mathbb{W}_3$, respectively. By the same token, $\Pi : \mathbb{W}_1$ indicates the proportion of stress that contributes to dissipative heating via (5.5.1), while $\Pi : \mathbb{W}_3$ indicates the proportion of perpendicular stress that contributes to transverse but non-dissipative transport of momentum.

The results of this analysis are plotted in Figure 5.11. The first row plots contours of the norm of the shear stress tensor, $|\mathbb{W}|$, in units of a reference shear frequency which we define as $\omega_s = v_{ti}/\alpha$, the ratio of the thermal velocity to initial interface width. As the vortex evolves, the magnitude and complexity of the velocity shear structures increase, presenting increased difficulty for leading-order gyroviscous stress closures.

The second row of Figure 5.11 plots $(\Pi : \mathbb{W}_1[\boldsymbol{u}])/(p_{ref}\omega_s)$. Notably, this quantity exhibits no discernable bias in one direction or another and is quite small, remaining less than 1% for the entire simulation lifetime. This indicates that systematic errors in the stress tensor closure for this simulation do not omit substantial amounts of dissipative viscous stress. The third row of the figure indicates the opposite conclusion for the component of Π in the direction of $\mathbb{W}_3[\boldsymbol{u}]$, which does exhibit a persistent bias in the positive direction. This indicates that the leading-order gyroviscous stress closure systematically underestimates the stress, despite the inclusion of the heat flux correction term. The final row of Figure 5.11 plots the same quantity for the Braginskii gyroviscous stress. Comparing the third and fourth rows, we conclude that while the leading-order gyroviscous stress is an improvement over Braginskii and substantially reduces the underestimation error, it does not eliminate it.

5.6 Discussion

We have evaluated the performance of the drift-ordered moment closures of Chapter 3 on a family of simulations which exhibit temperature gradients, sheared flow, and vorticity. The results indicate good agreement between the leading-order transport closures and kinetic moments, well into the nonlinear regime. The agreement, as measured by the R^2 value of the transport closure model, is tightly correlated with magnetization as measured by $\omega_c\tau$. On the other hand, agreement between closures and kinetic moments is lower when gradient length scales are near the ion Larmor radius as occurs during late-time turbulent dynamics.

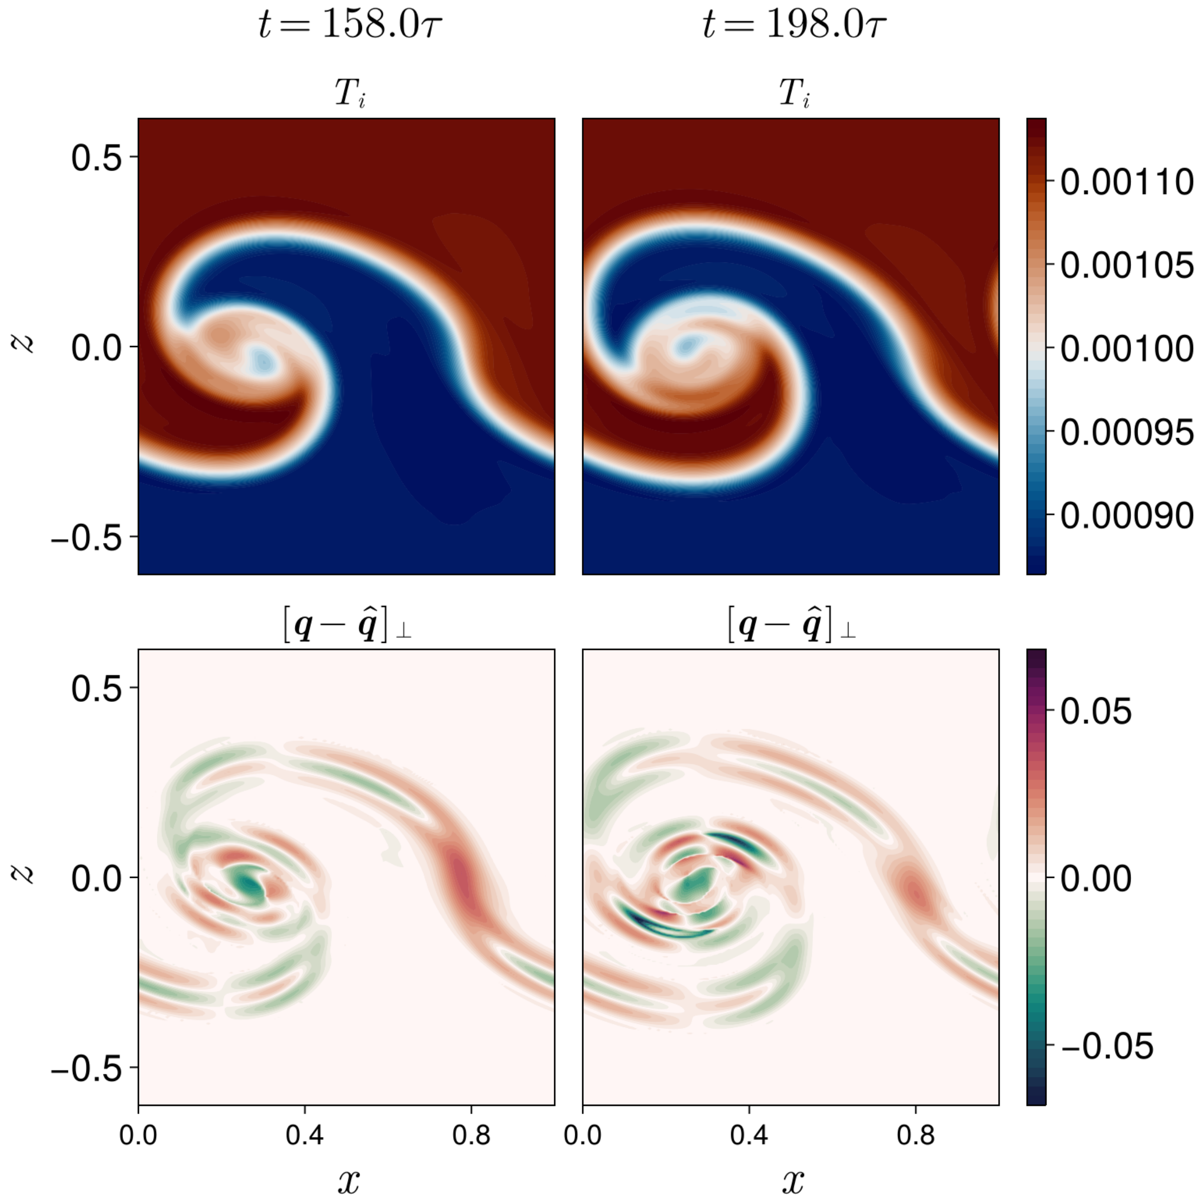

Figure 5.10: Contour plots of the normalized perpendicular component of the residual $\boldsymbol{q} - \hat{\boldsymbol{q}}$ for case S1. **First row**: contours of ion temperature. **Second row**: contours of $[\boldsymbol{q} - \hat{\boldsymbol{q}}]_\perp$, defined as the component of $\boldsymbol{q} - \hat{\boldsymbol{q}}$ in the direction of ∇T. Contours are plotted in units of the reference free-streaming heat flux $\boldsymbol{q}_{fs} = v_{ti} p_{ref}$. Dependence of sign of diffusion on the slope of the vortex matches the dependence observed in simulation M3 (compare to the third row of Figure 5.9.)

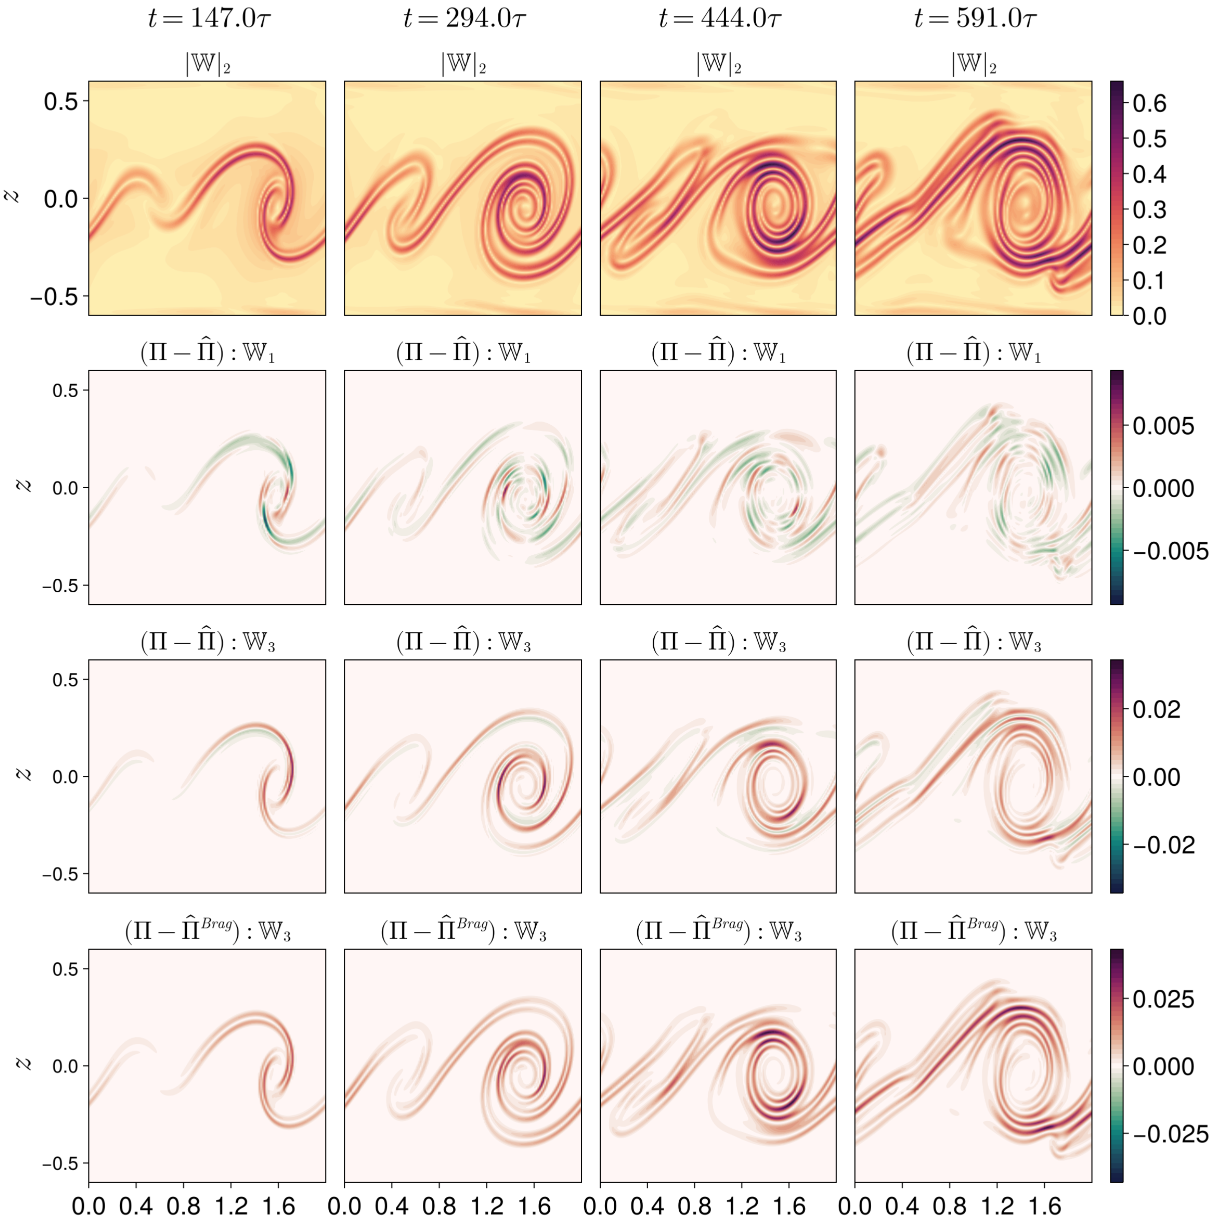

Figure 5.11: Components of the stress tensor residual from simulation M3, with $\omega_c\tau = 3.0$. **First row:** L^2 norm of $\nabla \boldsymbol{u}$ in units of v_{ti}/α. **Second row:** the magnitude of the component of $\Pi - \hat{\Pi}$ in the direction of $\mathbb{W}_1[\boldsymbol{u}]$. **Third row:** magnitude of the component of $\Pi - \hat{\Pi}$ in the direction of $\mathbb{W}_3[\boldsymbol{u}]$. **Fourth row:** magnitude of the Braginskii gyroviscous stress residual $\Pi - \hat{\Pi}^{Brag}$ in the direction of $\mathbb{W}_3[\boldsymbol{u}]$. Note that colorbars in rows 2-4 are centered at 0 to facilitate interpretation of signed quantities. Rows 2 and 3 show no bias in the direction of $\mathbb{W}_1$, but a significant bias in the direction of $\mathbb{W}_3$, indicating that gyroviscous transport of transverse momentum is greater in kinetic solutions than predicted by leading-order closures. Row 4 illustrates the systematic underestimation of kinetic stress by the Braginskii closure.

We have also examined the residuals of the diamagnetic heat flux and gyroviscous stress tensor transport closures. The heat flux residuals indicate clear dependence on higher-order derivatives of temperature. The heat flux residuals also include large components parallel to the temperature gradient, indicating that diffusive and anti-diffusive transport of heat may play a role in sheared-flow vortex physics. For the residuals of the gyroviscous stress, we have performed the analysis in the tensor space spanned by $\mathbb{W}_1$ and $\mathbb{W}_3$, which correspond to dissipative viscous and non-dissipative gyroviscous stress respectively. We determined that the higher-order residual terms do not contain an appreciable component in the direction of dissipative viscous stress. On the other hand, the gyroviscous stress is systematically underestimated by the leading-order closure in the simulation we examined, albeit by less than the classical Braginskii closure.

Finally, we proposed and evaluated an adjustment to the Braginskii gyroviscous stress tensor closure which is designed to be more accurate and feasible for implementers. The adjustment attempts to account for contributions to gyroviscous stress from the heat flux, using the diamagnetic heat flux closure and an estimate of the magnitude of the temperature gradient as an affordable approximation. The adjustment is most important under two conditions: first, when the diamagnetic drift makes up a substantial portion of ion velocity compared to the $E \times B$ drift. Second, when temperature gradients make a substantial contribution to the pressure gradients driving that diamagnetic drift. The size of the adjustment is summarized in the factor

$$\gamma^{Brag} = \frac{|\nabla p/p + \nabla T/T|}{|\nabla p/p|},$$

which is applied to the portion of gyroviscous stress due to the diamagnetic drift. The classical Braginskii closure underestimates or overestimates gyroviscous stress when $\gamma^{Brag} > 1$ or $\gamma^{Brag} < 1$, respectively.

The adjusted Braginskii gyroviscous stress is an approximation to the leading-order heat flux-corrected gyroviscous stress (5.3.1). Thus, the formal validity of the adjustment is independent of magnetization, and should be applied whenever the slow-dynamics ordering assumptions described in Chapter 3 hold. This necessitates an estimate of temperature gradient magnitude for use in γ^{Brag}. In this study we made use of global estimates of γ^{Brag}, which we found to be quite reliable. However, it may be the case that an accurate, numerically affordable estimate of γ^{Brag} is not available. Our numerical studies indicate that, while it may overestimate or underestimate the magnitude of gyroviscosity, the classical Braginskii closure is directionally correct in the cases considered here and can be used if the proposed adjustment is not feasible.

The adjusted Braginskii closure is of interest for computational modeling of gyroviscous processes in both multi-species [81, 98] and finite Larmor radius magnetohydrodynamic models [63]. In an MHD context the ion diamagnetic drift velocity can be estimated by assuming, for example, that all current density is diamagnetic. Then the same comments about estimating γ^{Brag} from either global considerations or local considerations apply.

We end this chapter by commenting on a pair of numerical studies [107, 102] of the magnetized Kelvin-Helmholtz instability and the kinetic stress tensors arising therein. Both studies find the Braginskii gyroviscous stress closure to be a good, but not perfect, predictor of the ion stress tensor. The simulations in [102] initialize a uniform plasma thermal pressure, and a non-uniform density profile, and thus have an infinite value of γ^{Brag}, meaning that a finite portion of the gyroviscous stress is completely missed by the Braginskii closure. In [107], a nearly isothermal kinetic equilibrium is initialized, but at late times in the simulation the ion temperature is found to vary significantly as a result of non-adiabaticity. Both studies use relatively large $E \times B$ velocities compared to the values simulated here. These studies of collisionless ion stress in the fast-dynamics regime urge us to consider exploring, in future work, the applicability of the heat flux correction and our affordable adjustment to these Kelvin-Helmholtz instabilities.

Chapter 6

Dynamical low-rank approximation

Up until now we have considered classical model reduction techniques for kinetic equations based on physical and asymptotic arguments. We have seen that such approaches, which typically result in fluid models for plasma, can successfully incorporate some elements of kinetic physics. However, we have also seen that fluid models can fall short when the plasma dynamics depart from their regime of asymptotic validity. In the remainder of this thesis, we will focus on a promising, non-asymptotic model reduction technique for kinetic equations: dynamical low-rank (DLR) approximation. DLR methods can, we will show, complement a traditional moment-based fluid approach and shore up some of their shortcomings in kinetics-dominated regimes.

This chapter is dedicated to a presentation of the DLR method in an abstract formulation. In chapters 7 and 8 we present a pair of extensions to the standard formulation which allow the method to obtain more physically accurate solutions to the Vlasov equation and related equations of plasma physics.

DLR approximation is best introduced as a family of methods for the solution of very large matrix-valued ordinary differential equations (ODEs) [71]. Many partial differential equations (PDEs) from physics can be naturally formulated as a matrix-valued ODE via the Method of Lines, and the resulting matrices are often very large indeed. Once equipped with the correct form of the algorithm, however, it becomes clear that one does not need to discretize first—the DLR algorithm is just as cogent when formulated as a decomposition of functions of two continuous variables, or two semidiscrete, vector-valued functions of one continuous variable.

To begin our discussion, we consider a large system of ordinary differential equations, whose unknowns may be arranged into an $m \times n$ time-dependent matrix $A(t)$. Such a system may be written simply as

$$\frac{d}{dt}A(t) = h(A, t). \tag{6.0.1}$$

If the dimensions of A are very large, the direct numerical solution of this system may be prohibitively expensive. Dynamical low-rank approximation refers to a class of techniques for solving systems of the form (6.0.1) efficiently, by representing the solution as a low-rank matrix. In particular, the solution is evolved directly in the low-rank format, without ever needing to form the full matrix A except to analyze the final solution. This can result in substantial computational cost savings depending on the numerical rank of the solution.

In this chapter we review low-rank approximation of matrices, and introduce the key ideas of dynamical low-rank approximation in the matrix-valued setting. This includes the critically important "robust integrators", which can stably integrate the resulting low-rank ODEs even in the presence of vanishing singular values, a necessary condition for convergence of the low-rank approximation. We then demonstrate how DLR approximation may be applied to partial differential equations directly without first discretizing, and even to semi-discrete equations with both discrete and continuous dimensions.

6.1 Basics of dynamical low-rank approximation

6.1.1 Low-rank approximation via the Singular Value Decomposition

The compression technique employed by DLR approximation is a low-rank approximation, where a large matrix $A \in \mathbb{R}^{m\times n}$ is approximated by the sum of r rank-1 matrices:

$$A \approx \tilde{A} = \sum_{i,j=1}^{r} \mathbf{u_i} \sigma_{\mathbf{ij}} \mathbf{v_j^T} \tag{6.1.1}$$

Since a rank one matrix is the outer product of vectors, we need store only the "skeleton" vectors $\mathbf{u}_i$ and $\mathbf{v}_j$, each of which require $\mathcal{O}(m)$ and $\mathcal{O}(n)$ storage, respectively. The total storage required for the approximation $\tilde{A}$ is therefore $\mathcal{O}(rm + rn + r^2)$, counting the storage required for the $r \times r$ matrix of coefficients, σ_{ij}.

This format is closely related to the compact Singular Value Decomposition (SVD), which gives for any rank-$\tilde{r}$ matrix $A \in \mathbb{R}^{m\times n}$ a factorization

$$A = U\Sigma V^T = \sum_{i}^{\tilde{r}} \mathbf{u}_i \sigma_i \mathbf{v}_i^T . \tag{6.1.2}$$

Here, the matrices $U \in \mathbb{R}^{m\times\tilde{r}}$ and $V \in \mathbb{R}^{n\times\tilde{r}}$ have orthonormal columns, which are called the left and right singular vectors. The matrix $\Sigma \in \mathbb{R}^{\tilde{r}\times\tilde{r}}$ is diagonal with non-negative entries, and typically has its entries arranged in descending order. With that ordering, the matrices U and V are unique up to multiplication of pairs of corresponding columns by -1.

It is well-known that truncations of the SVD provide the best rank-r approximation to a matrix [99]. To whit, define the truncated SVD of rank r as follows:

Definition 6.1.1 (Truncated SVD)**.** Let $A \in \mathbb{R}^{m\times n}$ have rank $\tilde{r} \leq \min\{m, n\}$, and let

$$A = \sum_{i}^{\tilde{r}} \mathbf{u}_i \sigma_i \mathbf{v}_i^T \tag{6.1.3}$$

be a compact Singular Value Decomposition of A, with the singular values arranged in descending order, $\sigma_1 \geq \sigma_2 \geq \ldots \sigma_{\tilde{r}}$. Then the truncated SVD with rank $r \leq \tilde{r}$ is

$$\tilde{A}_r = \sum_{i}^{r} \mathbf{u}_i \sigma_i \mathbf{v}_i^T = U_r \Sigma_r V_r^T \tag{6.1.4}$$

It is clear that the factors U_r and V_r^T also have orthonormal columns. We can make the connection to (6.1.1) by noting that we may perform any QR decomposition of both U_r and V_r, leading to the form

$$\tilde{A}_r = Q_U R_U \Sigma_r R_V^T Q_V^T. \tag{6.1.5}$$

The orthonormal columns of the factors Q_U and Q_V take the place of $\mathbf{u}_i$ and $\mathbf{v}_i$ in (6.1.1). The matrix $S_r = R_U \Sigma_r R_V^T$ is not diagonal, but it has the same singular values as $\tilde{A}_r$. The approximation given in (6.1.1) can therefore be seen as a variant of the SVD with a possibly non-diagonal matrix of singular values.

6.1.2 Dynamical low-rank approximation of matrices

We have seen that the truncated SVD provides the best rank-r approximation to a given matrix. When dealing with time-dependent problems, however, what is really wanted is a *trajectory* in the space of rank-r matrices which closely approximates the evolution of a time-dependent matrix $A(t)$. If we can exhibit an equation with a low-rank solution, whose dynamics are a good approximation to the dynamics of (6.0.1), then by integrating the low-rank equation directly, we can hopefully obtain a good approximation to $A(t)$. The benefit of this dynamic approach is that it never requires forming a matrix of size $m \times n$.

Formally, suppose we are given a matrix-valued ODE of the form of (6.0.1) with right-hand side $h(A, t)$. We would like to approximate its solution in the manifold $\mathcal{M}_r$ of rank-r matrices. A general member of $\mathcal{M}_r$ may be written in the form (6.1.1), or as a product of three matrix factors U, S, V^T as follows:

$$\mathcal{M}_r = \left\{ Y \in \mathbb{R}^{m\times n} :\ Y = USV^T,\ U^T U = I_r,\ V^T V = I_r,\ S \in \mathbb{R}^{r\times r} \text{ invertible} \right\}. \tag{6.1.6}$$

We write I_r to denote the identity matrix of dimension r. The dynamical low-rank approximation technique introduced in [71] proposes to integrate an approximation to the underlying ODE along $\mathcal{M}_r$,

$$\frac{d}{dt} Y(t) = \dot{Y}(t) \in \mathcal{T}_{Y(t)} \mathcal{M}_r \tag{6.1.7}$$

where $\mathcal{T}_{Y(t)}\mathcal{M}_r$ denotes the tangent space to $\mathcal{M}_r$ at $Y(t)$. It is natural to require a minimization of the form

$$\dot{Y}(t) = \text{argmin}_{\delta Y \in \mathcal{T}_{Y(t)}\mathcal{M}_r} \|\delta Y - h(Y(t), t)\|_F, \tag{6.1.8}$$

where $\|\cdot\|_F$ represents the Frobenius norm. That is, the time derivative of $\dot{Y}(t)$ is the trajectory which most closely approximates the dynamics h of the underlying equation evaluated at the current approximation $Y(t)$. Because the tangent space is a linear space, the variational principle (6.1.8) is equivalent to a Galerkin condition,

$$\left\langle \dot{Y}(t) - h(Y,t), \delta Y \right\rangle = 0, \quad \forall \delta Y \in \mathcal{T}_{Y(t)}\mathcal{M}_r. \tag{6.1.9}$$

The inner product $\langle \cdot \rangle$ should be understood in the sense of a Euclidean inner product between the vectors consisting of the stacked columns of each matrix $\dot{Y}, h, \delta Y$, whose corresponding norm is the Frobenius norm. Our goal is to translate this Galerkin condition into corresponding equations for each of the low-rank factors. Following [71], we consider the tangent space at U of the manifold $\mathcal{V}_{m,r}$ of matrices with orthonormal columns of dimension $m \times r$. This tangent space has the form

$$\mathcal{T}_U \mathcal{V}_{m,r} = \{\delta U \in \mathbb{R}^{m\times r} : \delta U^T U + U^T \delta U = 0\} = \{\delta U \in \mathbb{R}^{m \times r} : U^T \delta U \in \text{so}(r)\}. \tag{6.1.10}$$

The relation $\delta U^T U + U^T \delta U = 0$ may be justified by observing that we must have

$$\begin{aligned} I_m = (U + \delta U)^T (U + \delta U) &= U^T U + \delta U^T U + U^T \delta U + \cancel{\delta U^T \delta U} \\ &= I_m + \delta U^T U + U^T \delta U. \end{aligned}$$

The space so(r) denotes the space of real skew-symmetric matrices, which has dimension $(r^2 - r)/2$ due to the constraints on the entries imposed by skew-symmetry. Similarly, we have a tangent space at V of the manifold $\mathcal{V}_{n,r}$ of matrices with orthonormal columns of dimension $n \times r$:

$$\mathcal{T}_V \mathcal{V}_{n,r} = \{\delta V \in \mathbb{R}^{n\times r} : \delta V^T V + V^T \delta V = 0\} = \{\delta V \in \mathbb{R}^{n \times r} : V^T \delta V \in \text{so}(r)\}. \tag{6.1.11}$$

Finally, we let $\delta S \in \mathbb{R}^{r \times r}$ be arbitrary. These definitions allow us to define a map from the tensor product of the tangent spaces of the low rank factors, to the tangent space of the products, as follows:

$$\begin{aligned} &\mathbb{R}^{r\times r} \times \mathcal{T}_U \mathcal{V}_{m,r} \times \mathcal{T}_V \mathcal{V}_{n,r} \to \mathcal{T}_Y \mathcal{M}_r \times \text{so}(r) \times \text{so}(r) : \\ &\qquad (\delta S, \delta U, \delta V) \mapsto (\delta U S V^T + U \delta S V^T + U S \delta V^T, U^T \delta U, V^T \delta V). \end{aligned} \tag{6.1.12}$$

Of course we need to verify that the range of this map is as we say. We have $U^T \delta U \in \text{so}(r)$ because all elements of $T_U \mathcal{V}_{m,r}$ satisfy that relation per (6.1.10). Likewise $V^T \delta V \in \text{so}(r)$ by (6.1.11). Finally, the first component of the range lies in $\mathcal{T}_Y\mathcal{M}_r$ by an expansion of the definition (6.1.6):

$$\begin{aligned}(U+\delta U)(S+\delta S)(V^T+\delta V^T) &= USV^T + \delta USV^T + U\delta SV^T + US\delta V^T + \mathcal{O}(\delta^2)\\ &= Y + \delta Y\end{aligned}$$

We now count the dimensions on either side of (6.1.12). Observe that elements of $\mathcal{T}_U\mathcal{V}_{m,r}$ have mr degrees of freedom, but are subject to a total of $(r^2+r)/2$ constraints of the form

$$\delta\mathbf{u}_i^T\mathbf{u}_j + \mathbf{u}_i^T\delta\mathbf{u}_j = 0.$$

The same holds for $\mathcal{T}_V\mathcal{V}_{n,r}$, leading to a total dimension of the domain of (6.1.12) of

$$r^2 + mr - \frac{r^2+r}{2} + nr - \frac{r^2+r}{2} = mr + nr - r.$$

On the right hand side, the dimension of $\mathcal{T}_Y\mathcal{M}_r$ can be obtained in a similar way, by counting the degrees of freedom $mr+nr+r^2$, and subtracting the number of constraints, namely r^2 for each of the orthogonality relationships, to obtain $mr+nr-r^2$. Thus, the total dimension of the codomain of (6.1.12) is

$$mr + nr - r^2 + \frac{r^2-r}{2} + \frac{r^2-r}{2} = mr + nr - r.$$

Therefore, the dimensions of the domain and codomain of (6.1.12) match. It is easy to see, too, that the null space of (6.1.12) is zero, since the fixed factors U, V have orthonormal columns, and S is nonsingular. Thus, (6.1.12) is a bijection. This shows that all elements of $\mathcal{T}_Y$ have the form

$$\delta Y = \delta USV^T + U\delta SV^T + US\delta V^T. \tag{6.1.13}$$

Furthermore, by fixing a value for δU^TU and δV^TV, we will be able to invert the mapping and obtain an expression for the increments in the low rank components, $\delta U, \delta S, \delta V$, given an increment in the low-rank product δY. This is referred to in [71] as a gauge condition, since in principle the choice is arbitrary. It is most convenient to choose

$$\delta U^TU = \delta V^TV = 0. \tag{6.1.14}$$

We are now ready to use the Galerkin condition (6.1.9) to obtain evolution equations for the low-rank factors.

Theorem 6.1.1 ([71], Proposition 2.1)**.** *Let $Y = USV^T$ be a member of the manifold M_r defined in* (6.1.6), *and let* $h_Y = h(Y(t), t)$ *be fixed. The unique* $\dot{Y} \in \mathcal{T}_Y\mathcal{M}_r$ *which satisfies*

$$\left\langle \dot{Y} - h_Y, \delta Y \right\rangle = 0, \quad \forall\, \delta Y \in \mathcal{T}_Y\mathcal{M}_r \tag{6.1.15}$$

is given by

$$\dot{Y} = \dot{U}SV^T + U\dot{S}V^T + US\dot{V}^T, \tag{6.1.16}$$

where:

- $\dot{U} \in \mathcal{T}_U\mathcal{V}_{m,r}$ *is given by*

$$\dot{U} = (I - UU^T)h_Y VS^{-1}, \tag{6.1.17}$$

- $\dot{V} \in \mathcal{T}_V\mathcal{V}_{n,r}$ *is given by*

$$\dot{V} = (I - VV^T)h_Y^T US^{-T}, \tag{6.1.18}$$

- $\dot{S} \in \mathbb{R}^{r\times r}$ *is given by*

$$\dot{S} = U^T h_Y V. \tag{6.1.19}$$

Proof. We note that we have already shown that $\dot{Y}$ may be written in the form (6.1.16), which differs from (6.1.13) only in notation. Begin by letting $\delta Y = \mathbf{u}_i\mathbf{v}_j^T$. This is of the form (6.1.13) for $\delta U = \delta V = 0$ and $\delta S = \delta_{ij}$. Plugging this into (6.1.15) gives

$$\left\langle \dot{U}SV^T + U\dot{S}V^T + US\dot{V}^T - h_Y, \mathbf{u}_i\mathbf{v}_j^T \right\rangle = 0.$$

Recalling that the inner product $\langle\cdot\rangle$ is meant element-wise, this may be expanded to

$$\mathbf{u}_i^T(\dot{U}SV^T + U\dot{S}V^T + US\dot{V}^T)\mathbf{v}_j - \mathbf{u}_i^T h_Y \mathbf{v}_j = 0.$$

Using the gauge conditions $U^T\dot{U} = V^T\dot{V} = 0$, the first and third terms vanish. On the other hand, the orthogonality relations $U^TU = V^TV = I_r$ simplify the second term, leaving

$$\dot{s}_{ij} = \mathbf{u}_i^T h_Y \mathbf{v}_j.$$

This holds for each element of $\dot{S}$, from which we infer (6.1.19).

Now, let $\delta\mathbf{u}$ be arbitrary with $U^T\delta\mathbf{u} = 0$, and let $\delta Y = \sum_{j=1}^{r} \delta\mathbf{u} s_{ij}\mathbf{v}_j^T$. Then the Galerkin condition becomes

$$\delta\mathbf{u}^T(\dot{U}SV^T + U\dot{S}V^T + US\dot{V}^T)(VS^T)_i - \delta\mathbf{u}^T h_Y (VS^T)_i = 0,$$

where by $(VS^T)_i$ we denote the ith column of VS^T. This holds for all i, so we have, after cancellations,

$$\delta\mathbf{u}^T(\dot{U}S - h_Y V)S^T = 0.$$

By assumption S is invertible so we may right-multiply by S^{-T}:

$$\delta\mathbf{u}^T(\dot{U}S - h_Y V) = 0.$$

Since $\delta\mathbf{u}^T$ is an arbitrary vector orthogonal to the column space of U, this states that $\dot{U}S - h_Y V$ is in the null space of the orthogonal projector $P_U^{\perp} = I - UU^T$:

$$P_U^{\perp}\dot{U}S = P_U^{\perp} h_Y V.$$

However, since $\dot{U}$ is already orthogonal to U, $P_U^{\perp}\dot{U} = \dot{U}$, so we have

$$\dot{U} = (I - UU^T)h_Y V S^{-1}$$

as desired.

Finally, letting $\delta\mathbf{v}$ be arbitrary with $V^T\delta\mathbf{v} = 0$ and $\delta Y = \sum_{i=1}^{r} \mathbf{u}_i s_{ij}\delta\mathbf{v}_j^T$ allows us to derive (6.1.18) by a symmetrical argument. ☐

We have shown how to derive the equations of motion for the low rank factors. This leads to an algorithm for the dynamical low-rank approximation of a matrix equation of the form (6.0.1). Given a differential equation $\frac{dA}{dt} = h(A,t)$, and an initial condition $A(0) = A_0 \in \mathbb{R}^{m\times n}$, the dynamical low-rank algorithm computes an approximation $Y(t) \approx A(t)$ as follows:

- Fix a rank $r \leq \min\{m, n\}$
- Compute initial low-rank factors $U_0 S_0 V_0^T = \tilde{A}_{0_r}$ using the Truncated SVD (Definition 6.1.1).
- Solve the system of equations

$$\begin{cases} \frac{dU}{dt} = (I - UU^T)h(Y(t))VS^{-1}, & U(0) = U_0, \\ \frac{dV}{dt} = (I - VV^T)h(Y(t))^T US^{-T}, & V(0) = V_0, \\ \frac{dS}{dt} = U^T h(Y(t))V, & S(0) = S_0. \end{cases} \tag{6.1.20}$$

- As desired, reconstruct the approximate solution via

$$Y(t) = U(t)S(t)V^T(t). \tag{6.1.21}$$

Note that this algorithm seems to require computing $h(Y(t))$ as part of the right hand side of the low-rank evolution equations. In fact, the algorithm will only be efficient if it is not required to form the full matrix $Y(t)$ and the products $U^T h(USV^T)V$ can be reduced to simpler calculations. This happens when the function h has some tensor-friendly structure that can be exploited, such as in the Vlasov equation as we will see.

6.2 Robust dynamical low-rank integrators

The equations (6.1.17) and (6.1.18) require inverting the matrix S. In the case where the singular values of S become vanishingly small, S will be poorly conditioned, and its inversion is subject to extreme numerical instability. In theory it is still possible to follow the above algorithm, but the timesteps required for numerically stable integration may become impractically small. Unfortunately, the occurrence of vanishing singular values coincides with the approximation $Y = USV^T$ being a good one! There is therefore an upper limit on the quality of the approximation, imposed by the ill-conditioning of the algorithm just described. This motivates the introduction and use of a pair of dynamical low-rank integrators which are robust to vanishing singular values of S, which is sometimes called *overapproximation.* There are two such integrators: the "projector-splitting" integrator (PSI), and the Basis-Update & Galerkin (BUG) integrator. In this section we introduce both of them.

6.2.1 The projector-splitting integrator

The first robust integrator, introduced in [79], was the projector-splitting integrator. It is based on the idea that the Galerkin condition for the low-rank evolution $\dot{Y}$, given in (6.1.15), is an orthogonal projection onto the tangent space $\mathcal{T}_Y\mathcal{M}_r$. In fact, this projection can be written in an explicit form, as the following theorem shows:

Theorem 6.2.1. *The projection $P(Y)$ defined by the Galerkin condition*

$$\langle P(Y)h_Y - h_Y, \delta Y\rangle = 0, \quad \forall\, \delta Y \in \mathcal{T}_Y\mathcal{M}_r$$

has the explicit form

$$P(Y)h_Y = h_Y VV^T - UU^T h_Y VV^T + UU^T h_Y. \tag{6.2.1}$$

Proof. The proof follows directly from plugging in the equations (6.1.17), (6.1.18), (6.1.19) into (6.1.16):

$$\begin{aligned}
P(Y)h_Y = \dot{Y} &= (I - UU^T)h_Y V S^{-1} S V^T + UU^T h_Y V V^T + USS^{-1} U^T h_Y (I - VV^T) \\
&= (I - UU^T)h_Y VV^T + UU^T h_Y VV^T + UU^T h_Y (I - VV^T) \\
&= h_Y VV^T - UU^T h_Y VV^T + UU^T h_Y
\end{aligned}$$

□

Thus, we can rewrite the dynamical low-rank approximation equation as

$$\begin{aligned}
\frac{dY}{dt} &= P(Y)h(Y,t) \\
&= h(Y,t)VV^T - UU^T h(Y,t)VV^T + UU^T h(Y,t).
\end{aligned}$$

This form lends itself immediately to a first-order operator splitting method:

$$\begin{cases}
\frac{dY_1}{dt} = h(Y_1,t)V_1V_1^T & t \in [t_0, t_0+\Delta t],\ Y_1(t_0) = Y(t_0) \\
\frac{dY_2}{dt} = -U_2U_2^T h(Y_2,t)V_2V_2^T & t \in [t_0, t_0+\Delta t],\ Y_2(t_0) = Y_1(t_0+\Delta t) \\
\frac{dY_3}{dt} = U_3U_3^T h(Y_3,t) & t \in [t_0, t_0+\Delta t],\ Y_3(t_0) = Y_2(t_0+\Delta t) \\
Y(t_0+\Delta t) = Y_3(t_0+\Delta t).
\end{cases} \tag{6.2.2}$$

Subscripts on the factors U and V here correspond to the low-rank factorization of the corresponding stage, i.e. $Y_2 = U_2 S_2 V_2^T$. Higher-order splitting schemes are also possible, using standard operator-splitting techniques. However, in this form the algorithm is not yet complete, because we are not interested in integrating Y in time directly. The splitting needs to correspond to an algorithm that operates at the level of the low-rank factors U, S, and V. Luckily, it does correspond to such an algorithm, which is due to the fact that each of the sub-steps stays of rank r. The following theorem gives the practical projector-splitting algorithm which operates at the level of the low-rank factors.

Theorem 6.2.2 ([79], Lemma 3.1)**.** *The solution of* (6.2.2) *is equivalent to the solution of the following scheme:*

- $Y_1(t)$ *is given by* $Y_1(t) = U_1(t)S_1(t)V_1(t)$*, where the* U_1, S_1, V_1 *are the solutions to*

$$\begin{aligned}
&\frac{d(U_1S_1)}{dt} = h(Y_1,t)V_1, \quad \frac{dV_1}{dt} = 0, \quad t \in [t_0, t_0+\Delta t] \\
&U_1(t_0) = U(t_0),\ S_1(t_0) = S(t_0),\ V_1(t_0) = V(t_0)
\end{aligned}$$

- *$Y_2(t)$ is given by $Y_2(t) = U_2(t)S_2(t)V_2^T$, where U_2, S_2, V_2 are the solutions to*

$$\begin{aligned}&\frac{dS_2}{dt} = -U_2^T h(Y_2, t)V_2, \quad \frac{dU_2}{dt} = \frac{dV_2}{dt} = 0, \quad t \in [t_0, t_0 + \Delta t],\\ &U_2(t_0) = U_1(t_0 + \Delta t),\ S_2(t_0) = S_1(t_0 + \Delta t),\ V_2 = V_1(t_0 + \Delta t)\end{aligned}$$

- *$Y_3(t)$ is given by $Y_3(t) = U_3(t)S_3(t)V_3(t)^T$, where U_3, S_3, V_3 are the solutions to*

$$\begin{aligned}&\frac{d(S_3V_3^T)}{dt} = U_3^T h(Y_3, t), \quad \frac{dU_3}{dt} = 0, \quad t \in [t_0, t_0 + \Delta t]\\ &U_3(t_0) = U_2(t_0 + \Delta t),\ S_3(t_0) = S_2(t_0 + \Delta t), V_3(t_0) = V_2(t_0 + \Delta t).\end{aligned}$$

Proof. Observe that the right hand side of each of the substeps of (6.2.2) is in the tangent space $T_Y\mathcal{M}_r$ due to the orthogonality of the low-rank factors. For example,

$$\begin{aligned}P(Y_1)[h(Y_1, t)V_1V_1^T] &= h(Y_1, t)V_1V_1^TV_1V_1^T - U_1U_1^Th(Y_1, t)V_1V_1^TV_1V_1^T + U_1U_1^Th(Y_1, t)V_1V_1^T\\ &= h(Y_1, t)V_1V_1^T.\end{aligned}$$

Therefore, the solution $Y_1(t)$ of (6.2.2) stays of rank r for the whole time interval $[t_0, t_0 + \Delta t]$, and may be factorized as

$$Y_1(t) = U_1(t)S_1(t)V_1(t)^T,$$

with each of the low-rank factors satisfying the usual criteria. Plugging this into (6.2.2) and applying a product rule yields

$$\frac{dY_1}{dt} = \frac{d(U_1S_1)}{dt}V_1^T + U_1S_1\frac{dV_1^T}{dt} = h(Y_1, t)V_1V_1^T.$$

It is clear that taking $\frac{d(U_1S_1)}{dt} = h(Y_1, t)V_1$ and $\frac{dV_1}{dt} = 0$ satisfies this equation.

The proof for the other two steps is similar. □

Theorem 6.2.2 leads to an algorithm for the efficient solution of ODEs in the low-rank format. The final ingredient is the application of the QR decomposition to extract the low-rank factors U, S from their product US after the first step, and vice versa for the factors S, V^T from their product after the third step. Here we introduce notation which we will use throughout the rest of our presentation, and which is standard in the literature, namely the "auxiliary" bases K and L.

Algorithm 6.2.1. Given a right hand side $h(Y, t)$ and low-rank factors $U(t_0), S(t_0)$, and $V(t_0)$ at time t_0, compute the factors $U(t_1), S(t_1)$, and $V(t_1)$ at time $t_1 = t_0 + \Delta t$ in the following way:

- **K step**:
  - Solve the initial value problem

    $$\frac{dK}{dt} = h(K(t)V(t_0)^T, t)V(t_0), \quad K(t_0) = U(t_0)S(t_0) \tag{6.2.3}$$

    for $K(t_1)$.
  - Perform a QR decomposition to factorize $K(t_1) = U(t_1)S'$.
- **S step**:
  - Solve the initial value problem

    $$\frac{dS''}{dt} = U(t_1)^T h(U(t_1)S''(t)V(t_0)^T, t)V(t_0), \quad S''(t_0) = S' \tag{6.2.4}$$

    for $S''(t_1)$.
- **L step**:
  - Solve the initial value problem

    $$\frac{dL}{dt} = U(t_1)^T h(U(t_1)L(t), t), \quad L(t_0) = S''(t_1)V(t_0) \tag{6.2.5}$$

    for $L(t_1)$.
  - Perform a QR decomposition to factorize $L(t_1) = S(t_1)V(t_1)$.

The Basis Update & Galerkin Integrator

Another popular robust integrator for DLR systems is the Basis Update & Galerkin (BUG) integrator, which was introduced in [16]. We note a pair of advantages of the BUG integrator over the projector-splitting integrator [16]:

- The BUG integrator has a pair of substeps which correspond to the K and L steps of the PSI. Whereas in the PSI there is a data dependency between these substeps, the BUG integrator can perform them in parallel.
- The PSI contains a backwards time step for the S step, which may be problematic for PDEs with diffusive behavior for example: a backwards step is not always well-posed. The BUG integrator takes a forward S step. As indicated in [19] and [2], this may not be as large a problem for the PSI as it seems, even for diffusive problems, since a Forward Euler step for the anti-diffusive term is stable in practice.

The numerical methods in this thesis use the PSI, so we will not derive the BUG integrator in detail. The high-level algorithm is as follows:

Algorithm 6.2.2. Given a right hand side $h(Y,t)$ and low-rank factors $U(t_0), S(t_0)$, and $V(t_0)$ at time t_0, compute the factors $U(t_1), S(t_1)$, and $V(t_1)$ at time $t_1 = t_0 + \Delta t$ in the following way:

- **Parallel K & L steps**:
  - Solve the initial value problem

    $$\frac{dK}{dt} = h(K(t)V(t_0)^T, t)V(t_0), \quad K(t_0) = U(t_0)S(t_0) \tag{6.2.6}$$

    for $K(t_1)$.
  - Perform a QR decomposition to factorize $K(t_1) = U(t_1)R_1$ and compute the matrix $M = U(t_1)^T U(t_0)$.
  - Solve the initial value problem

    $$\frac{dL}{dt} = U(t_1)^T h(U(t_1)L(t), t), \quad L(t_0) = V(t_0)(S''(t_1))^T \tag{6.2.7}$$

    for $L(t_1)$.
  - Perform a QR decomposition to factorize $L(t_1) = V(t_1)\tilde{R}_1^T$ and compute the matrix $N = V(t_1)^T V(t_0)$.
- **S step**: Solve the initial value problem

  $$\frac{dS}{dt} = U(t_1)^T h(U(t_1)S(t)V^T(t_1))V(t_1), \quad S(t_0) = MS(t_0)N^T \tag{6.2.8}$$

  for $S(t_1)$.

6.3 Dynamical low-rank approximation of kinetic equations

In this section we describe the application of the dynamical low-rank integrators of Section 6.1.2 to kinetic equations posed on a 1D1V dimensional phase space. As usual, our notation will take x, or $\boldsymbol{x}$ in the multidimensional case, as the coordinate parametrizing physical space, and v or $\boldsymbol{v}$ as the coordinate parametrizing velocity space. A prototypical kinetic equation is

$$\partial_t f + \boldsymbol{v} \cdot \nabla_{\boldsymbol{x}} f + \boldsymbol{F} \cdot \nabla_{\boldsymbol{v}} f = C(f).$$

We take this equation to be posed over a Cartesian product domain,

$$(\boldsymbol{x}, \boldsymbol{v}, t) \in \Omega_x \times \mathbb{R}^d \times \mathbb{R}^+, \tag{6.3.1}$$

where $\Omega_x \subset \mathbb{R}^d$ is any physical domain, while the velocity is allowed to range over all of $\mathbb{R}^d$.

In this work we consider decompositions of f into $\boldsymbol{x}$-dependent and $\boldsymbol{v}$-dependent factors. Other decompositions are also possible, for example higher-order tensor decompositions which further separate the different velocity coordinates; see for example [36]. The basic low-rank decomposition we consider is

$$f = \sum_{ij} X_i(\boldsymbol{x}, t) S_{ij}(t) V_j(\boldsymbol{v}, t). \tag{6.3.2}$$

As usual, sums over indices i, j run from 1 to r. We note that this decomposition already departs from the form of the matrix decomposition (6.1.1) since its factors are continuous functions of the independent variables $\boldsymbol{x}$ and $\boldsymbol{v}$. The development of the dynamical low-rank approximation method proceeds in much the same way for functions as for matrices, so we will omit the full derivation. Instead, we will point out the key elements. First, we require the spatial, velocity, and phase-spatial domains to be equipped with L^2 inner products, which we denote with angle brackets:

$$\begin{aligned} \langle X_i, X_j \rangle_x &= \int_{\Omega_x} X_i(\boldsymbol{x}) X_j(\boldsymbol{x}) \, \mathrm{d}\boldsymbol{x}, \quad \langle V_i, V_j \rangle_v = \int_{\mathbb{R}^d} V_i(\boldsymbol{v}) V_j(\boldsymbol{v}) \, \mathrm{d}\boldsymbol{v}, \\ \langle f, g \rangle_{xv} &= \int_{\Omega_x} \int_{\mathbb{R}^d} f(\boldsymbol{x}, \boldsymbol{v}) g(\boldsymbol{x}, \boldsymbol{v}) \, \mathrm{d}\boldsymbol{v} \, \mathrm{d}\boldsymbol{x}. \end{aligned} \tag{6.3.3}$$

The requirement of orthogonality of the basis functions is written in terms of these inner products, i.e.

$$\langle X_i, X_j \rangle_x = \delta_{ij}, \quad \langle V_i, V_j \rangle_v = \delta_{ij} \tag{6.3.4}$$

The low-rank manifold $\mathcal{M}_r$ will consist of functions which may be written in the form (6.3.2), with the requirement of orthogonality of the basis vectors:

$$\begin{aligned} \mathcal{M}_r = \Bigg\{ & f : \Omega_x \times \mathbb{R}^d \to \mathbb{R}^+ : f(\boldsymbol{x}, \boldsymbol{v}) = \sum_{ij} X_i(\boldsymbol{x}) S_{ij} V_j(\boldsymbol{v}), \\ & \qquad \langle X_i, X_j \rangle_x = \delta_{ij}, \quad \langle V_i, V_j \rangle_v = \delta_{ij}, \quad S \in \mathbb{R}^{r \times r} \text{ invertible} \Big\}. \end{aligned}$$

The corresponding tangent space is then defined in an analogous manner to (6.1.10). The DLR time step is determined by a Galerkin projection in the phase space inner product, namely

$$\left\langle \dot{f}(t) - h(f, t), \delta f \right\rangle_{xv} = 0, \quad \forall \delta f \in \mathcal{T}_f \mathcal{M}_r. \tag{6.3.5}$$

Here, as before, $h(f,t)$ denotes the entire right-hand side of the evolution equation for f, as in $\partial_t f = h(f,t)$. Finally, the projector corresponding to this Galerkin condition may be written in terms of projections along each component:

$$P(f)h = \sum_j \langle h, V_j \rangle_v V_j - \sum_{ij} X_i \langle X_i V_j, h \rangle_{xv} V_j + \sum_i X_i \langle X_i, h \rangle_x . \tag{6.3.6}$$

With this form of the projector we can construct direct analogues of the DLR Algorithm 6.2.1. If the form of h is amenable, then the resulting algorithm will be asymptotically much faster than any algorithm for solving the "full tensor" kinetic equation.

Indeed, kinetic equations are particularly appealing targets for the dynamical low-rank approximation technique for two complementary reasons. The first is that much of the difficulty of solving them comes from their high dimensionality, which DLR approximation is specifically suited to address. The second reason is that the terms in kinetic equations tend to be "tensor-friendly" in the sense that they are simple products of space-dependent and velocity-dependent factors. Substituting a low-rank ansatz such as (6.3.2) into the free transport term, for example, gives a term which remains low-rank:

$$\boldsymbol{v} \cdot \nabla_{\boldsymbol{x}} f = \boldsymbol{v} \cdot \nabla_{\boldsymbol{x}} \sum_{ij} X_i(\boldsymbol{x},t) S_{ij}(t) V_j(\boldsymbol{v},t) = \sum_{ij} S_{ij} [\nabla_{\boldsymbol{x}} X_i(\boldsymbol{x},t)] \cdot [\boldsymbol{v} V_j(\boldsymbol{v},t)]. \tag{6.3.7}$$

This property is essential for efficient implementation of the DLR integrators, which all involve the projection of such a low-rank representation onto the subspaces spanned by the low-rank bases. Only terms which are separable or low-rank in this sense can be efficiently projected.

Remark 1 (Notation)**.** When treating multidimensional kinetic equations in the low-rank framework, we must be careful that our notation distinguishes between the different kinds of inner products that we will encounter. These are:

- Inner products of size d. We will denote these with a dot, such as $\boldsymbol{v} \cdot \nabla_{\boldsymbol{x}}$. Vectors with dimensional components will be denoted by boldface letters.
- Inner products of size r: Sums over the ranks of the low-rank factors. These will usually be written out explicitly as sums over indices i, j, k, l.
- Inner products of size N: Products between two vectors or a matrix and a vector, as when applying a flux Jacobian matrix to a vector of unknowns in quasilinear form. We will denote column vectors by upper case boldface letters or lowercase letters such as $\boldsymbol{u}, \boldsymbol{q}$. Matrices will be denoted by uppercase non-bolded letters. An example of this sort of notation is

$$\boldsymbol{W}_j^T A \boldsymbol{W}_l, \tag{6.3.8}$$

representing a quadratic form applied to the vectors $\boldsymbol{W}_j$ and $\boldsymbol{W}_l$.

- Inner products of continuous dimension: function inner products are denoted by angle brackets, and are presumed to be in the usual L^2 inner product space over the domain of definition of the functions. For example,

$$\langle X_i, \partial_x X_k \rangle_x = \int_{\Omega_x} X_i(x) \partial_x X_k(x) \, \mathrm{d}x. \tag{6.3.9}$$

Example 6.3.1. To illustrate the approach, we present the DLR treatment of the collisionless Vlasov equation in one dimension,

$$\partial_t f + v \partial_x f + E \partial_v f = 0, \tag{6.3.10}$$

which is coupled with Gauss's law,

$$E = -\partial_x \phi, \quad -\partial_x^2 \phi = \rho, \quad \rho = \int_{\mathbb{R}} f \, \mathrm{d}v.$$

The following derivation is due to [36]. Begin by isolating the time derivative on the left hand side to define h:

$$\partial_t f = -v \partial_x f - E \partial_v f \triangleq h.$$

We now plug h and the low-rank ansatz (6.3.2) into the substeps of the projector-splitting integrator of Algorithm 6.2.1, to obtain

- **K step**: Integrate the following equation with initial condition $K_j^n = \sum_i X_i^n S_{ij}^n$ to obtain K_j^{n+1}:

$$\partial_t K_j = \langle V_j, h \rangle_v = -\sum_l \langle V_j, v V_l \rangle_v \partial_x K_l - \sum_l E \langle V_j, \partial_v V_l \rangle_v K_l, \tag{6.3.11}$$

- Perform a QR decomposition to obtain X^{n+1}, S':

$$K_j^{n+1} = \sum_i X_i^{n+1} S'_{ij}.$$

- **S step:** Integrate the following equation with initial condition S' to obtain S'':

$$\partial_t S_{ij} = -\langle X_i V_j, h \rangle_{xv} = \sum_{kl} \langle X_i, \partial_x X_k \rangle_x \langle V_j, v V_l \rangle_{xv} + \sum_{kl} \langle X_i, E X_k \rangle_x \langle V_j, \partial_v V_l \rangle_v . \tag{6.3.12}$$

- **L step:** Integrate the following equation with initial condition $L_i^n = \sum_j S''_{ij} V_j^n$ to obtain L_i^{n+1}:

$$\partial_t L_i = \langle X_i, h \rangle_x = -\sum_k \langle X_i, \partial_x X_k \rangle_x v L_k - \sum_k \langle X_i, E X_k \rangle_x \partial_v L_k. \tag{6.3.13}$$

- Perform a QR decomposition to obtain S^{n+1}, V^{n+1}:

$$L_i^{n+1} = \sum_j S_{ij}^{n+1} V_j^{n+1}.$$

The above steps give a first-order in time algorithm for the dynamical low-rank approximate solution of (2.1.3). It must be combined with appropriate spatial and temporal discretizations of the subflows (6.3.11), (6.3.12), and (6.3.13), to give a fully discrete algorithm.

It is worth pointing out the form of the equations obtained from the projector-splitting DLR algorithm. From a scalar-valued equation with a pair of advection terms, (2.1.3), we obtained two systems of PDEs for K and L respectively. Each is an r-dimensional hyperbolic source-balance law. For the K flow, the flux Jacobian is

$$\langle V_j, v V_l \rangle_v ,$$

and for the L flow

$$\langle X_i, E X_k \rangle_x .$$

In both cases, the flux Jacobian is symmetric, and thus has a set of r real distinct eigenvalues, making the system strictly hyperbolic.

Chapter 7

An efficient DLR algorithm for the Vlasov-Ampère-Fokker-Planck equation

This chapter describes a dynamical low-rank (DLR) algorithm for solving a plasma kinetic equation with strong collisions. The strongly collisional regime is one of the two important asymptotic limits we discussed in Chapter 1. Because of the multiscale nature of many plasmas, we are interested in developing algorithms that can capture kinetic physics, while also *efficiently* capturing fluid physics in regions where the fluid description is valid.

One idea along these lines is hybridization [59, 24], which applies a kinetic description to regions of the plasma that are anticipated to require it, and an efficient fluid description elsewhere. Hybridization can greatly increase simulation efficiency, but requires foreknowledge of the plasma dynamics to decompose the domain effectively into fluid and kinetic subdomains. Moreover, it can introduce solution artifacts at subdomain boundaries if the kinetic solution there is not sufficiently close to a Maxwellian.

The algorithm described in this chapter approaches the question from a different angle. Rather than using a fluid model to capture the behavior of near-Maxwellian distribution functions, we use a kinetic description throughout the solution domain. Our contribution is a DLR algorithm which is specifically designed to become more efficient as the solution approaches a local Maxwellian, by requiring lower rank in that regime.

The work covered here has been published in [19]. The chapter is organized as follows: In Section 7.1 we present the model equation considered in this chapter, which we call the Vlasov-Ampère-Fokker-Planck (VAFP) equation, and its fluid limit. Section 7.2 presents a dynamical low-rank approximation method for the VAFP equation and estimates the computational cost associated with each substep. Section 7.3 describes how the semi-discrete system of DLR equations is made fully discrete by the introduction of discretizations in physical and velocity space. In section 7.4 we show that, in the small-ϵ limit, the semi-discrete system is a consistent discretization of the limiting fluid system. This result is a formalization of the motivation for this method: it shows that we capture the fluid limit efficiently. Section 7.5 presents numerical validation of the proposed discrete scheme on a test problems. Section 7.6 closes the chapter with remarks on how the scheme handles the anti-diffusive S step.

7.1 The Vlasov-Ampère-Fokker-Planck system and its fluid limit

The kinetic equation we consider in this chapter is the linearized Vlasov-Fokker-Planck equation:

$$\partial_t f + \boldsymbol{v} \cdot \nabla_{\boldsymbol{x}} f + \frac{1}{\epsilon} \boldsymbol{E} \cdot \nabla_{\boldsymbol{v}} f = \frac{1}{\epsilon} \nabla_{\boldsymbol{v}} \cdot (\boldsymbol{v} f + \nabla_{\boldsymbol{v}} f). \tag{7.1.1}$$

This is coupled with Ampère's law for the electric field,

$$\partial_t \boldsymbol{E} = -\boldsymbol{J}, \quad \boldsymbol{J}(x,t) = \langle \boldsymbol{v}, f \rangle_{\boldsymbol{v}}\,. \tag{7.1.2}$$

Here we have made use of a notation for the L^2 inner product, which we define as

$$\langle g, h \rangle_{\boldsymbol{x}} = \int_{\Omega} gh \, \mathrm{d}\boldsymbol{x}, \quad \langle g, h \rangle_{\boldsymbol{v}} = \int_{\mathbb{R}^d} gh \, \mathrm{d}\boldsymbol{v}, \quad \langle g, h \rangle_{\boldsymbol{x},\boldsymbol{v}} = \int_{\Omega} \int_{\mathbb{R}^d} gh \, \mathrm{d}\boldsymbol{x} \, \mathrm{d}\boldsymbol{v}.$$

Equation (7.1.1) is a significant simplification of the full Vlasov-Maxwell-Landau model given by equations (2.1.3) and (2.1.5), but it retains key properties of that equation. Namely, (7.1.1) models a diffusive collision operator and nonlinear coupling between f and $\boldsymbol{E}$, both of which are key difficulties to solving the Vlasov-Maxwell-Landau system.

The scaling represented by the small parameter ϵ has been called the "high-field" limit of the Vlasov equation. It was introduced in [85] as a model for the semiconductor Boltzmann equation, was treated numerically in [15], and in [67] with an asymptotic-preserving scheme.

Because (7.1.1) is a single-species kinetic equation, which describes the motion of an electron species against a static ion background, we must supply a corresponding static background charge density, which we denote $\eta(\boldsymbol{x})$. This represents the density of ions, which do not move on the timescale resolved by (7.1.1). The field $\boldsymbol{E}$ should satisfy Gauss's law,

$$\boldsymbol{E} = -\nabla_{\boldsymbol{x}} \phi(\boldsymbol{x}), \quad -\nabla_{\boldsymbol{x}}^2 \phi(\boldsymbol{x}) = \rho(\boldsymbol{x}, t) - \eta(\boldsymbol{x}), \tag{7.1.3}$$

where

$$\rho(\boldsymbol{x}, t) = \langle 1, f \rangle_{\boldsymbol{v}} .$$

To recapitulate, in this chapter we solve the coupled system

$$\begin{cases} \partial_t f + \boldsymbol{v} \cdot \nabla_{\boldsymbol{x}} f + \frac{1}{\epsilon} \boldsymbol{E} \cdot \nabla_{\boldsymbol{v}} f = \frac{1}{\epsilon} \nabla_{\boldsymbol{v}} \cdot (\boldsymbol{v} f + \nabla_{\boldsymbol{v}} f), & (x, v, t) \in \Omega \times \mathbb{R}^d \times [0, T] \\ \partial_t \boldsymbol{E} = -\boldsymbol{J}, & (x, t) \in \Omega \times [0, T] \\ E = -\nabla_{\boldsymbol{x}} \phi, \quad -\nabla_{\boldsymbol{x}}^2 \phi = \rho - \eta & x \in \Omega, t = 0. \end{cases} \tag{7.1.4}$$

Asymptotic fluid limit

The limit of (7.1.4) for small ϵ is an electrostatic "creeping flow", in which inertial forces are vanishingly small compared to electrostatic forces. To analyze the limit $\epsilon \to 0$, we introduce a scaled "local Maxwellian" defined as

$$M(\boldsymbol{x}, \boldsymbol{v}, t) = \frac{1}{(2\pi)^{d/2}} e^{-\frac{|\boldsymbol{v} - \boldsymbol{E}(\boldsymbol{x}, t)|^2}{2}}. \tag{7.1.5}$$

In contrast to the Maxwellian equilibrium of the Boltzmann equation, this function has a uniform density. It is isothermal and its flow velocity is equal to $\boldsymbol{E}$. It is not hard to show [19] that (7.1.1) is equivalent to

$$\partial_t f + \boldsymbol{v} \cdot \nabla_{\boldsymbol{x}} f = \frac{1}{\epsilon} \nabla_{\boldsymbol{x}} \cdot \left[M \nabla_{\boldsymbol{v}} (M^{-1} f) \right]. \tag{7.1.6}$$

The form (7.1.6) reveals the dominant balance of (7.1.1). As $\epsilon \to 0$, formally $\nabla_{\boldsymbol{v}} \cdot \left[M \nabla_{\boldsymbol{v}} (M^{-1} f) \right] \to 0$, which implies $f \to \rho M$. Details of this derivation can be found in [19].

We can derive a macroscopic system in the limit as $\epsilon \to 0$ by taking moments with respect to 1 and $\boldsymbol{v}$ of (7.1.6), to obtain

$$\partial_t \rho + \nabla_{\boldsymbol{x}} \cdot \boldsymbol{J} = 0, \tag{7.1.7}$$

$$\partial_t J + \nabla_{\boldsymbol{x}} \cdot \langle \boldsymbol{v} \otimes \boldsymbol{v}, f \rangle_{\boldsymbol{v}} = \frac{1}{\epsilon} (\rho \boldsymbol{E} - \boldsymbol{J}). \tag{7.1.8}$$

Formally, from (7.1.8), we have $\boldsymbol{J} \to \rho \boldsymbol{E}$ as $\epsilon \to 0$. In this limit, (7.1.7) becomes

$$\partial_t \rho + \nabla_{\boldsymbol{x}} \cdot (\rho \boldsymbol{E}) = 0, \tag{7.1.9}$$

which when combined with Ampère's law gives the limiting fluid system

$$\begin{cases}\partial_t \rho + \nabla_{\boldsymbol{x}} \cdot (\rho \boldsymbol{E}) = 0, & (x,t) \in \Omega \times [0,T] \\ \partial_t \boldsymbol{E} = -\rho \boldsymbol{E}, & (x,t) \in \Omega \times [0,T] \\ E = -\nabla_{\boldsymbol{x}} \phi, \quad -\nabla_{\boldsymbol{x}}^2 \phi = \rho - \eta & x \in \Omega, t = 0. \end{cases} \tag{7.1.10}$$

7.2 Dynamical low-rank algorithm

We reiterate that our goal is to derive a DLR algorithm for the solution of (7.1.4) which is efficient near the fluid limit (7.1.10). As we saw, in the limit as $\epsilon \to 0$, f approaches ρM. This limit is not a low-rank function of $\boldsymbol{x}$ and $\boldsymbol{v}$, because M is not (7.1.5). However, the function $f M^{-1} = \rho(\boldsymbol{x}, t)$ clearly is a low-rank function of $\boldsymbol{x}$ and $\boldsymbol{v}$: it is independent of $\boldsymbol{v}$.

We therefore propose a low-rank ansatz, not for f, but for the quotient $g = f/M$, which has rank one in the fluid limit:

$$g(\boldsymbol{x}, \boldsymbol{v}, t) \approx \sum_{ij}^{r} X_i(\boldsymbol{x}, t) S_{ij}(t) V_j(\boldsymbol{v}, t). \tag{7.2.1}$$

Here, r is the rank of the low-rank representation.

The governing equation for g can be found by substitution, and is

$$\partial_t g = -\boldsymbol{v} \cdot \nabla_{\boldsymbol{x}} g + \frac{1}{\epsilon}[(\nabla_{\boldsymbol{v}} - \boldsymbol{v} + \boldsymbol{E}) \cdot \nabla_{\boldsymbol{v}} g] - \mathcal{M} g \tag{7.2.2}$$

$$\triangleq h, \tag{7.2.3}$$

where

$$\mathcal{M} = \frac{1}{M}(\partial_t M + \boldsymbol{v} \cdot \nabla_{\boldsymbol{x}} M). \tag{7.2.4}$$

We solve (7.2.2) using the projector-splitting integrator (PSI). Recall that the PSI makes use of auxiliary bases K and L, defined as

$$K_j(\boldsymbol{x}, t) = \sum_i X_i(\boldsymbol{x}, t) S_{ij}(t), \quad L_i(\boldsymbol{v}, t) = \sum_j S_{ij}(t) V_j(\boldsymbol{v}, t).$$

Substitute (7.2.2) into the projection formula (6.3.6), and splitting into the K, S, and L steps in the usual way gives the following system of equations:

$$\begin{aligned}
\partial_t K_j &= -\sum_l (\nabla_{\boldsymbol{x}} K_l) \cdot \langle \boldsymbol{v} V_j V_l \mathcal{M} \rangle_{\boldsymbol{v}} + \frac{1}{\epsilon} \sum_l K_l \left[\langle V_j (\nabla_{\boldsymbol{v}} - \boldsymbol{v}) \cdot \nabla_{\boldsymbol{v}} V_l \rangle_{\boldsymbol{v}} + \boldsymbol{E} \cdot \langle V_j \nabla_{\boldsymbol{v}} V_l \rangle_{\boldsymbol{v}} \right], && (7.2.5)\\
\partial_t S_{ij} &= \sum_{kl} S_{kl} \langle X_i \nabla_{\boldsymbol{x}} X_k \rangle_{\boldsymbol{x}} \cdot \langle \boldsymbol{v} V_j V_l \rangle_{\boldsymbol{v}} + \sum_{kl} S_{kl} \langle X_i X_k V_j V_l \mathcal{M} \rangle_{\boldsymbol{x}\boldsymbol{v}}\\
&\quad - \frac{1}{\epsilon} \sum_{kl} S_{kl} \left[\langle X_i X_k \rangle_{\boldsymbol{x}} \langle V_j (\nabla_{\boldsymbol{v}} - \boldsymbol{v}) \cdot \nabla_{\boldsymbol{v}} V_l \rangle_{\boldsymbol{v}} + \langle X_i X_k \boldsymbol{E} \rangle_{\boldsymbol{x}} \cdot \langle V_j \nabla_{\boldsymbol{v}} V_l \rangle_{\boldsymbol{v}} \right], && (7.2.6)\\
\partial_t L_i &= -\sum_k \boldsymbol{v} \cdot \langle X_i (\nabla_{\boldsymbol{x}} X_k) \rangle_{\boldsymbol{x}} L_k - \sum_k \langle X_i X_k \mathcal{M} \rangle_{\boldsymbol{x}} L_k + \frac{1}{\epsilon} \sum_k \left[\langle X_i X_k \rangle_{\boldsymbol{x}} (\nabla_{\boldsymbol{v}} - \boldsymbol{v}) + \langle X_i X_k \boldsymbol{E} \rangle_{\boldsymbol{x}} \right] \cdot \nabla_{\boldsymbol{v}} L_k. && (7.2.7)
\end{aligned}$$

We also expand the terms involving $\mathcal{M}$ (defined in (7.2.4)):

$$\partial_t M = (v - E) \cdot (\partial_t E) M = -M (v - E) \cdot J,$$

$$v \cdot \nabla_x M = M \sum_{i,j} (v_j - E_j) v_i \partial_{x_i} E_j = M \sum_{i,j} (v_j v_i \partial_{x_i} E_j - E_j v_i \partial_{x_i} E_j) = M (v \otimes v) : \nabla_x E - \frac{M}{2} v \cdot \nabla_x (E^2).$$

Note that $\frac{1}{2} \nabla_x (E^2) \neq E \cdot \nabla_x E$. $(\nabla_x E)_{ij} := \partial_{x_j} E_i$ and $A : B := \sum_{ij} a_{ij} b_{ij}$. Putting these together we obtain

$$\begin{aligned}
\mathcal{M} &= \frac{1}{M} (\partial_t M + v \cdot \nabla_x M) = E \cdot J - v \cdot J - \frac{1}{2} v \cdot \nabla_x (E^2) + (v \otimes v) : \nabla_x E\\
&:= \mathcal{M}_1 + v \cdot \mathbf{M}_2 + (v \otimes v) : \overline{\overline{\mathsf{M}}}_3, && (7.2.8)
\end{aligned}$$

where

$$\mathcal{M}_1 = E \cdot J, \quad \mathbf{M}_2 = -J - \frac{1}{2} \nabla_x (E^2), \quad \overline{\overline{\mathsf{M}}}_3 = \nabla_x E. \tag{7.2.9}$$

Here we use **boldface** to denote vectors of length d, and sans-serif to denote tensors of size $d \times d$. Both vectors and tensors may also vary in x and v. In all cases the tensor contractions $\cdot, :$ indicate contraction over the length-d dimensions. The terms involving $\mathcal{M}$ then expand to

$$\begin{aligned}
\langle V_j V_l \mathcal{M} \rangle_v &= \delta_{jl} \mathcal{M}_1 + \langle v V_j V_l \rangle_v \cdot \mathbf{M}_2 + \langle (v \otimes v) V_j V_l \rangle_v : \overline{\overline{\mathsf{M}}}_3,\\
\langle X_i X_k \mathcal{M} \rangle_x &= \langle X_i X_k \mathcal{M}_1 \rangle_x + v \cdot \langle X_i X_k \mathbf{M}_2 \rangle_x + (v \otimes v) : \langle X_i X_k \overline{\overline{\mathsf{M}}}_3 \rangle_x,\\
\langle X_i X_k V_j V_l \mathcal{M} \rangle_{xv} &= \delta_{jl} \langle X_i X_k \mathcal{M}_1 \rangle_x + \langle v V_j V_l \rangle_v \cdot \langle X_i X_k \mathbf{M}_2 \rangle_x\\
&\quad + \langle (v \otimes v) V_j V_l \rangle_v : \langle X_i X_k \overline{\overline{\mathsf{M}}}_3 \rangle_x.
\end{aligned}$$

The projector splitting scheme for $\tilde{g}$ may be straightforwardly coupled with Ampère's equation (7.1.2), which can be written in terms of the low-rank components as

$$\partial_t E = -\langle vM\tilde{g}\rangle_v = -\sum_{i,j} X_i S_{ij} \langle vMV_j\rangle_v. \tag{7.2.10}$$

7.2.1 Time-integration algorithm

The algorithm described up to this point has been fully continuous, except for the projection onto the low-rank manifold. We now present a discretization in time, leaving space continuous for the moment. The time discretization makes use of an implicit-explicit (IMEX) scheme for capturing the fast dynamics of the collision operator in the fluid limit ($\epsilon \ll 1$).

In the following we report rough estimates of the computational complexity of each substep. To avoid complicating the presentation unnecessarily, for these estimates we consider $d \sim 1$, so that we are free to ignore both the dimension and constant factors in our "big-O" notation.

Suppose we have the quantities $(E^n, X_i^n, V_j^n, S_{ij}^n)$ at timestep t^n. Then we calculate $(E^{n+1}, X_i^{n+1}, V_j^{n+1}, S_{ij}^{n+1})$ in the following way.

Update E

1. Compute the following integral appearing in (7.2.10):

$$I_j^n(x) := \langle vV_j^n M^n\rangle_v = \frac{1}{(2\pi)^{d/2}} \int vV_j^n(v) e^{-\frac{|v-E^n(x)|^2}{2}}\, \mathrm{d}\boldsymbol{v}. \tag{7.2.11}$$

   A naive computation of this integral requires $\mathcal{O}(N_x N_v)$ steps, a computational cost that is unacceptably high. However, since the Maxwellian is isothermal, the integral has a convolutional structure, and may be computed using a Fast Fourier Transform (FFT). The required substeps are:

   - Compute the convolution

$$\ell_j^n(\zeta) = \left[(v \mapsto vV_j^n(v)) * \left(v \mapsto e^{-|v|^2/2} \right) \right] (\zeta) \tag{7.2.12}$$

     using an FFT.

     Cost: $\mathcal{O}(rN_v \log N_v)$.

   - Compute the composition of E^n with ℓ_j^n using any interpolation scheme from the FFT nodes to an arbitrary point $E^n(x)$:

$$I_j^n(x) = \frac{1}{(2\pi)^{d/2}} \ell_j^n(E^n(x)).$$

     Cost: $\mathcal{O}(rN_x)$.

Exploiting the convolutional structure of $I_j^n(x)$ with an FFT reduces the total computational cost to $\mathcal{O}(rN_v \log N_v)$, which is acceptable.

2. Compute the current density:

$$J^n(x) = \sum_{ij} X_i^n(x) S_{ij}^n I_j^n(x). \tag{7.2.13}$$

Cost: $\mathcal{O}(r^2 N_x)$.

3. Perform a Forward Euler step to solve (7.2.10):

$$E^{n+1}(x) = E^n(x) - \Delta t J^n. \tag{7.2.14}$$

Cost: $\mathcal{O}(N_x)$.

Remark 2 (Solving Gauss's law)**.** If one chooses to solve the Vlasov-Poisson system instead of the Vlasov-Ampére system considered here, the E update requires the solution of an elliptic equation for the potential ϕ:

$$-\nabla_x^2 \phi^n = \rho^n - \eta, \quad E^{n+1} = -\nabla_x \phi^n. \tag{7.2.15}$$

The only aspect requiring special treatment in the low-rank method is the computation of the charge density ρ^n appearing on the right-hand side of Gauss's law. This can be computed in much the same manner as the current density J^n in equations (7.2.11)-(7.2.13). Briefly, compute r integrals

$$Y_j^n(x) := \left\langle V_j^n M^n \right\rangle_v = \frac{1}{(2\pi)^{d/2}} \int V_j^n(v) e^{-\frac{|v-E^n|^2}{2}} \, \mathrm{d}v. \tag{7.2.16}$$

This may be done by computing convolutions

$$y_j^n(\zeta) = \left[\left(v \mapsto V_j^n(v) \right) \star \left(v \mapsto e^{-|v|^2/2} \right) \right] (\zeta), \tag{7.2.17}$$

and interpolating in any desired manner to the arbitrary points $E^n(x)$:

$$Y_j^n(x) = \frac{1}{(2\pi)^{d/2}} y_j^n(E^n(x)). \tag{7.2.18}$$

Then the charge density is given by

$$\rho^n(x) = \sum_{ij} X_i^n(x) S_{ij}^n Y_j^n(x). \tag{7.2.19}$$

It will be seen that this procedure is nearly identical to that for computing J^n, with the only difference that one computes the zeroth moment of f rather than the first moment. Note that in this formulation, one still needs to calculate J^n via the corresponding convolutions, since it will be used to update K, S, and L as described below.

Remark 3. The use of the FFT to achieve a fast convolutional algorithm is similar to that proposed in [33]. Computationally the approach is nearly identical, with $E^n(x)$ taking the place of $u^n(x)$ as the local drift velocity of the Maxwellian.

K step

1. Compute integrals in v. We use **boldface** to denote vector-valued matrices of total size $r \times r \times d$, and sans-serif to denote tensor-valued matrices of total size $r \times r \times d \times d$. In both cases the indices running over the length-d dimensions are suppressed. The integrals to compute are:

$$\mathbf{c}^1_{jl} = \langle v V^n_j V^n_l \rangle_v, \quad \bar{\bar{\mathsf{c}}}^2_{jl} = \langle (v \otimes v) V^n_j V^n_l \rangle_v, \tag{7.2.20}$$

$$d^1_{jl} = \langle V^n_j (\nabla_v - v) \cdot \nabla_v V^n_l \rangle_v, \quad \mathbf{d}^2_{jl} = \langle V^n_j \nabla_v V^n_l \rangle_v. \tag{7.2.21}$$

   Cost: $\mathcal{O}(r^2 N_v)$.

2. Compute $\mathcal{M}^n_1, \mathbf{M}^n_2$, and $\bar{\bar{\mathsf{M}}}^n_3$:

$$\mathcal{M}^n_1 = E^n \cdot J^n, \quad \mathbf{M}^n_2 = -J^n - \frac{1}{2}\nabla_x((E^n)^2), \tag{7.2.22}$$

$$\bar{\bar{\mathsf{M}}}^n_3 = \nabla_x E^n. \tag{7.2.23}$$

   Cost: $\mathcal{O}(N_x)$.

3. Compute matrices on the right hand side of (7.2.5).

$$A^1_{jl} = \delta_{jl}\mathcal{M}^n_1 + \mathbf{c}^1_{jl} \cdot \mathbf{M}^n_2 + \bar{\bar{\mathsf{c}}}^2_{jl} : \bar{\bar{\mathsf{M}}}^n_3, \tag{7.2.24}$$

$$A^2_{jl} = d^1_{jl} + E^n \cdot \mathbf{d}^2_{jl}. \tag{7.2.25}$$

   Cost: $\mathcal{O}(r^2 N_x)$.

4. The evolution equation (7.2.5) for K may now be written as

$$\partial_t K_j = -\sum_l \mathbf{c}^1_{jl} \cdot \nabla_x K_l - \sum_l A^1_{jl} K_l + \frac{1}{\epsilon}\sum_l A^2_{jl} K_l. \tag{7.2.26}$$

Advance (7.2.26) in time, using an IMEX step to handle the stiff term:

$$\sum_l \left[\delta_{jl} - \frac{\Delta t}{\epsilon} A^2_{jl}\right] K^{n+1}_l = K^n_j - \Delta t \left(\sum_l \mathbf{c}^1_{jl} \cdot \nabla_x K^n_l + \sum_l A^1_{jl} K^n_l \right). \tag{7.2.27}$$

Note that the only differential operator, namely ∇_x, appearing in this equation is treated explicitly. Therefore the linear system appearing in this equation involves no coupling between points in x. When discretized it will consist of N_x separate systems each of size $r \times r$. We can solve this small system at each point in x using any standard dense linear solver—the size is not large enough to warrant any special technique.

Cost: $\mathcal{O}(r^2 N_x)$ for both the right-hand side and the implicit step, due to solving each $r \times r$ system separately.

5. Perform a QR decomposition of K^{n+1}_j to obtain X^{n+1}_i and S^1_{ij}.

   Cost: $\mathcal{O}(r^2 N_x)$.

S step

1. Compute the integrals in x, using the new basis X^{n+1}:

$$c^\star_{ik} = \langle X^{n+1}_i X^{n+1}_k \mathcal{M}^n_1 \rangle_x, \quad \mathbf{c}^{\star\star}_{ik} = \langle X^{n+1}_i X^{n+1}_k \mathbf{M}^n_2 \rangle_x, \quad \overline{\overline{\mathsf{c}}}^{\star\star\star}_{ik} = \langle X^{n+1}_i X^{n+1}_k \overline{\overline{\mathsf{M}}}^n_3 \rangle_x, \tag{7.2.28}$$

$$\mathbf{d}^\star_{ik} = \langle X^{n+1}_i \nabla_x X^{n+1}_k \rangle_x, \quad \mathbf{e}^\star_{ik} = \langle X^{n+1}_i X^{n+1}_k E^n \rangle_x. \tag{7.2.29}$$

   Cost: $\mathcal{O}(r^2 N_x)$.

2. Compute

$$\tilde{c}_{ij;kl} = \langle X^{n+1}_i X^{n+1}_k V^n_j V^n_l \mathcal{M}^n \rangle_{xv} = \delta_{jl} c^\star_{ik} + \mathbf{c}^1_{jl} \cdot \mathbf{c}^{\star\star}_{ik} + \overline{\overline{\mathsf{c}}}^2_{jl} : \overline{\overline{\mathsf{c}}}^{\star\star\star}_{ik}. \tag{7.2.30}$$

   Cost: $\mathcal{O}(r^4)$.

3. Compute the order-four tensors

$$B^1_{ij;kl} = d^\star_{ik} \cdot \mathbf{c}^1_{jl} + \tilde{c}_{ij;kl}, \tag{7.2.31}$$

$$B^2_{ij;kl} = \delta_{ik} d^1_{jl} + \mathbf{e}^\star_{ik} \cdot \mathbf{d}^2_{jl}. \tag{7.2.32}$$

   Cost: $\mathcal{O}(r^4)$.

4. The evolution equation for S may now be written as

$$\partial_t S_{ij} = \sum_{kl} B^1_{ij;kl} S_{kl} - \frac{1}{\epsilon} \sum_{kl} B^2_{ij;kl} S_{kl}. \tag{7.2.33}$$

Perform a Forward Euler step to advance $S^1_{ij} \to S^2_{ij}$:

$$S^2_{ij} = S^1_{ij} + \Delta t \sum_{kl} S^1_{kl} B^1_{ij;kl} - \frac{\Delta t}{\epsilon} \sum_{kl} S^1_{kl} B^2_{ij;kl}. \tag{7.2.34}$$

Our use of a Forward Euler step here differs from the presentation in [33], where an IMEX step was used to advance the S equation in the case of the BGK collision operator. Since the structure of the Fokker-Planck operator is more complicated than the BGK type, special care is needed. Heuristically, one can see that something different from the K and L steps may be required, simply because the S equation runs backwards in time. For a more detailed justification and a discussion of how the situation differs from [33], refer to Section 7.6.

Cost: $\mathcal{O}(r^4)$.

L step

1. Compute the $r \times r$ matrix

$$\hat{c}_{ik} = c^\star_{ik} + v \cdot \mathbf{c}^{\star\star}_{ik} + (v \otimes v) : \overline{\overline{\mathbf{c}}}^{\star\star\star}_{ik}. \tag{7.2.35}$$

Cost: $\mathcal{O}(r^2)$.

2. The equation (7.2.7) may now be written as

$$\partial_t L_i = -\sum_k v \cdot \mathbf{d}^\star_{ik} L_k - \sum_k \hat{c}_{ik} L_k + \frac{1}{\epsilon} \sum_k \left(\delta_{ik} \nabla_v^2 - \delta_{ik} v \cdot \nabla_v + \mathbf{e}^\star_{ik} \cdot \nabla_v \right) L_k. \tag{7.2.36}$$

Advance (7.2.36) using an IMEX step by solving the system

$$\sum_k \left[\delta_{ik} - \frac{\Delta t}{\epsilon} (\delta_{ik} \nabla_v^2 - \delta_{ik} v \cdot \nabla_v + \mathbf{e}^\star_{ik} \cdot \nabla_v) \right] L^{n+1}_k = L^n_i - \Delta t \sum_k (v \cdot \mathbf{d}^\star_{ik} + \hat{c}_{ik}) L^n_k. \tag{7.2.37}$$

Note that in contrast to (7.2.27), the left-hand side of this equation does involve differential operators in v, and so the linear system may be discretized by a fully coupled (but sparse) matrix of size $rN_v \times rN_v$. There are $\mathcal{O}(rN_v)$ non-empty entries. Using Krylov subspace methods lets us keep the total cost of solving this system on the order of $\mathcal{O}(rN_v)$, assuming the number of iterations does not grow unboundedly with r or N_v, which is what we observe in practice.

Cost: $\mathcal{O}(r^2 N_v)$.

3. Perform a QR decomposition of L_i^{n+1} to obtain V_j^{n+1} and S_{ij}^{n+1}.

   Cost: $\mathcal{O}(r^2 N_v)$.

Adding together all of our computational complexity estimates, we get a total cost of $\mathcal{O}(r^4 + r^2 N_x + r^2 N_v)$ which compares quite favorably with the cost of the full tensor method $\mathcal{O}(N_x N_v)$.

7.3 Fully discrete algorithm

In this section we address the question of physical and velocity space discretization. One of the virtues of the dynamical low-rank method is that it decouples the discretization of the two bases, X_i and V_j, which may be treated more or less independently. The X basis is updated by solving a system of coupled hyperbolic PDEs in (7.2.27), while the V basis is updated by solving a parabolic system in (7.2.37). These systems are coupled via the matrix of singular values S, as well as weighted inner products of whichever basis is being held constant (viz. (7.2.20), (7.2.21), etc.) We are free to choose whichever discretization is most appropriate for the corresponding evolution equation of each basis. In this work we use second-order finite difference discretizations in both x and v for simplicity. In principle, it is easy to choose, for example, a Fourier spectral method to take advantage of periodicity in the x direction, or even a more involved method such as Discontinuous Galerkin along one or the other basis, without increasing the implementation complexity too greatly.

7.3.1 Spatial discretization

Our spatial discretization in x is designed to solve the explicit part of the evolution equation for K, which is (7.2.27). This is a linear hyperbolic PDE with the flux matrix $\mathbf{c}^1_{jl}$, which is a symmetric matrix. We opt for a second-order finite difference discretization with flux limiting, as described in [78], section 16.2. To illustrate, we consider the situation in two spatial dimensions, $d = 2$. The matrix $\mathbf{c}^1_{jl}$ consists of components $\mathbf{c}^{1;m}_{jl}$ for $m \in \{1, 2\}$, acting on the x and y directions respectively. The matrices $\mathbf{c}^{1;m}_{jl}$ are symmetric and real; recall their definition (7.2.20). Therefore they are unitarily diagonalizable, and we can write

$$(T^m)^T \mathbf{c}^{1;m} T^m = \sum_{jl} T^m_{ij} \mathbf{c}^{1;m}_{jl} T^m_{kl} = \lambda^m_i \delta_{ik} = \Lambda^m. \tag{7.3.1}$$

Left-multiplying (7.2.27) by $(T^1)^T$, and introducing the eigenbasis $\hat{K}^n_i = [(T^1)^T K^n]_i$, gives the system

$$\sum_{jl} T^1_{ij} \left[I - \frac{\Delta t}{\epsilon} \mathbf{A}^2_{jl} \right] K^{n+1}_l = \hat{K}^n_i - \Delta t \lambda^1_i \partial_x \hat{K}^n_i - \Delta t \sum_{jl} T^1_{ij} \left(\mathbf{c}^{1;2}_{jl} \partial_y K^n_l + \mathbf{A}^1_{jl} K^n_l \right). \tag{7.3.2}$$

At a grid point x_p, the flux-limited finite difference discretization approximates $\lambda_i^1 \partial_x \hat{K}_i^n(x_p)$ by a difference of fluxes at half grid points $x_{p+1/2}, x_{p-1/2}$:

$$\lambda_i^1 \partial_x \hat{K}_i^n(x_p) \approx \frac{F(\lambda_i, \hat{K}_i^n)_{p+1/2} - F(\lambda_i, \hat{K}_i^n)_{p-1/2}}{\Delta x}. \tag{7.3.3}$$

The flux F is given by the combination of a first-order flux (upwinding) flux,

$$F_L(\lambda, \hat{K})_{p+1/2} = \frac{\lambda}{2}(\hat{K}_{p+1} + \hat{K}_p) - \frac{|\lambda|}{2}(\hat{K}_{p+1} - \hat{K}_p), \tag{7.3.4}$$

with a second-order Lax-Wendroff flux. The combination is governed by a flux-limiter $\phi(\theta)$, which stabilizes the scheme in the presence of sharp changes in the gradient:

$$F(\lambda, \hat{K})_{p+1/2} = F_L(\lambda, \hat{K})_{p+1/2} + \frac{1}{2}\phi(\theta_{p+1/2}) \left(\operatorname{sgn}(\lambda) - \frac{\lambda \Delta t}{\Delta x} \right) \lambda \delta(\hat{K})_{p+1/2}, \tag{7.3.5}$$

where $\delta(\hat{K})_{p+1/2} = (\hat{K}_{p+1} - \hat{K}_p)$ The quantity $\theta_{p+1/2}$ measures how quickly the gradient is changing in the vicinity of $x_{p+1/2}$, and itself uses upwinding based on the sign of λ:

$$\theta_{p+1/2} = \frac{\delta(\hat{K})_{p+1/2-\operatorname{sgn}(\lambda)}}{\delta(\hat{K})_{p+1/2}}. \tag{7.3.6}$$

The function $\phi : \mathbb{R} \to [0, 2]$ is called the limiter, and there are many options to choose from. We use the Van Leer limiter,

$$\phi(\theta) = \frac{|\theta| + \theta}{1 + |\theta|}. \tag{7.3.7}$$

After approximating the term $\lambda_i^1 \partial_x \hat{K}_i^n$, for each eigenvalue λ_i, we can transform back to the original variables by left-multiplying with (T^1):

$$\sum_l \left[I - \frac{\Delta t}{\epsilon} \mathbf{A}_{jl}^2 \right] K_l^{n+1} = K_j^n - \Delta t \sum_i T_{ij}^1 \delta_x(\lambda_i, \hat{K}^n)_i - \Delta t \sum_l \left(\mathbf{c}_{jl}^{1;2} \partial_y K_l^n + \mathbf{A}_{jl}^1 K_l^n \right). \tag{7.3.8}$$

The discretization in y is handled similarly, by left-multiplying (7.3.8) by T^2. The above scheme is second-order in smooth regions of the solution, and degrades to first order around discontinuities and extrema.

7.3.2 Velocity discretization

Our discretization in v is designed to effectively solve (7.2.37), which is a parabolic system (strictly speaking, a convection-diffusion type equation). We recall the linear system to be solved here:

$$\sum_k \left[\delta_{ik} - \frac{\Delta t}{\epsilon} (\delta_{ik} \nabla_v^2 - \delta_{ik} v \cdot \nabla_v + \mathbf{e}_{ik}^\star \cdot \nabla_v) \right] L_k = RHS. \tag{7.3.9}$$

It is convenient to discretize this operator by splitting the left hand side into a diagonal ($i = k$) term and an offdiagonal term. In the case when $i = k$, we have

$$\begin{aligned}\left(\delta_{ik}\nabla_v^2 - \delta_{ik} v \cdot \nabla_v + \mathbf{e}^\star_{kk} \cdot \nabla_v\right) L_k &= \left(\nabla_v^2 - v \cdot \nabla_v + \mathbf{e}^\star_{kk} \cdot \nabla_v\right) L_k \\ &= \left[\left((\nabla_v - v) + \mathbf{e}^\star_{kk}\right) \cdot \nabla_v\right] L_k \\ &= \frac{1}{M^k} \nabla_v \cdot \left(M^k \nabla_v L_k\right) \\ &:= \mathbb{T}_k(\mathbf{e}^\star_{kk}) L_k, \end{aligned} \tag{7.3.10}$$

where M^k is the local Maxwellian

$$M^k = e^{-\frac{|v-\mathbf{e}^\star_{kk}|^2}{2}}. \tag{7.3.11}$$

In one dimension, a second-order-accurate central difference discretization of (7.3.10) is

$$(\mathbb{T}_k(\mathbf{e}^\star_{kk})L)_p \approx \frac{1}{M^k_p \Delta v}\left(M^k_{p+1/2}\frac{L_{p+1} - L_p}{\Delta v} - M^k_{p-1/2}\frac{L_{p-1} - L_p}{\Delta v}\right) \tag{7.3.12}$$

$$= \frac{M^k_{p+1/2}L_{p+1} - (M^k_{p+1/2} + M^k_{p-1/2})L_p + M^k_{p-1/2}L_{p-1}}{M^k_p \Delta v^2}. \tag{7.3.13}$$

The off-diagonal terms are simply

$$\mathbf{e}^*_{ik} \cdot \nabla_v L_k := \mathbb{U}_{ik} L_k. \tag{7.3.14}$$

For simplicity we discretize this using a second-order centered difference operator. Stability is not a concern, since it will be coupled to a Backwards Euler timestepping scheme. With these discretizations in hand the implicit step for L takes the form

$$\sum_k \left[\delta_{ik} - \frac{\Delta t}{\epsilon}\left(\delta_{ik}\mathbb{T}_k(\mathbf{e}^\star_{kk}) + \mathbb{U}_{ik}\right)\right] L_k^{n+1} = L_i^n - \Delta t \sum_k (v \cdot \mathbf{d}^\star_{ik} + \hat{c}_{ik}) L_k^n. \tag{7.3.15}$$

Despite being of size $rN_v \times rN_v$, this linear system is quite sparse, having roughly $\mathcal{O}(r^2 N_v)$ nonzero entries. It is therefore amenable to fast solution by iterative solvers. Since it is not symmetric, we use the Restarted GMRES [90] iterative algorithm. We find good results by preconditioning with the constant matrix $\left(\delta_{ik} - \delta_{ik}\frac{\Delta t}{\epsilon}\mathbb{T}_k(\mathbf{0})\right)^{-1}$. Timings of our code indicate that this step takes on the same order of magnitude as the other components of the algorithm, up to the largest problems we consider here.

7.4 Asymptotic behavior of the discrete scheme

In this section we demonstrate that the discrete scheme described in the preceding sections preserves the asymptotic limit (7.1.10) as $\epsilon \to 0$.

We consider the limit of the discrete system at the level of g, which is advanced via the K, S, and L steps with an accuracy that is first-order in time and second-order in space:

$$\frac{g^{n+1}-g^n}{\Delta t} = -v\cdot\nabla_x g^n - \mathcal{M}^n g^n + \frac{1}{\epsilon}(M^n)^{-1}\nabla_v\cdot(M^n\nabla_v g^{n+1}) + \mathcal{O}(\Delta t + \Delta x^2). \tag{7.4.1}$$

Furthermore, the electric field is advanced by

$$\frac{E^{n+1}-E^n}{\Delta t} = -J^n. \tag{7.4.2}$$

From (7.4.1), we can see that

$$(M^n)^{-1}\nabla_v\cdot(M^n\nabla_v g^{n+1}) = \mathcal{O}(\epsilon) \implies g^{n+1} = c + \mathcal{O}(\epsilon), \tag{7.4.3}$$

that is to say, after one time step, we expect the solution g to be close to a constant function in v when ϵ is small. To see this, just note the following

$$\int \nabla_v\cdot[M\nabla_v g]\log g\,\mathrm{d}\boldsymbol{v} = \int \nabla_v\cdot[Mg\nabla_v \log g]\log g\,\mathrm{d}\boldsymbol{v} = -\int Mg\left|\nabla_v \log g\right|^2\,\mathrm{d}\boldsymbol{v} \le 0, \tag{7.4.4}$$

where the equality holds if and only if g is a function independent of v. Moreover, (7.4.3) implies

$$\rho^{n+1} = \int M^{n+1}g^{n+1}\,\mathrm{d}\boldsymbol{v} = c\int M^{n+1}\,\mathrm{d}\boldsymbol{v} + \mathcal{O}(\epsilon) = c + \mathcal{O}(\epsilon), \tag{7.4.5}$$

$$J^{n+1} = \int vM^{n+1}g^{n+1}\,\mathrm{d}\boldsymbol{v} = c\int vM^{n+1}\,\mathrm{d}\boldsymbol{v} + \mathcal{O}(\epsilon) = cE^{n+1} + \mathcal{O}(\epsilon) = \rho^{n+1}E^{n+1} + \mathcal{O}(\epsilon). \tag{7.4.6}$$

On the other hand, we can multiply (7.4.1) by M^n and integrate in v to obtain

$$\begin{aligned}
\frac{\int M^n g^{n+1}\,\mathrm{d}\boldsymbol{v} - \rho^n}{\Delta t} &= -\int v\cdot(\nabla_x g^n)M^n\,\mathrm{d}\boldsymbol{v} - \int(\partial_t M^n + v\cdot\nabla_x M^n)g^n\,\mathrm{d}\boldsymbol{v} \\
&\quad + \frac{1}{\epsilon}\cancel{\int\nabla_v\cdot(M^n\nabla_v g^{n+1})}\,\mathrm{d}\boldsymbol{v} + \mathcal{O}(\Delta t + \Delta x^2) \\
&= -\int[v\cdot(\nabla_x g^n)M^n + (v\cdot\nabla_x M^n)g^n]\,\mathrm{d}\boldsymbol{v} - \partial_t E^n\cdot\int(v-E^n)M^n g^n\,\mathrm{d}\boldsymbol{v} \\
&\quad + \mathcal{O}(\Delta t + \Delta x^2) \\
&= -\nabla_x\cdot J^n - c\partial_t E^n\cdot\cancel{\int(v-E^n)M^n}\,\mathrm{d}\boldsymbol{v} + \mathcal{O}(\epsilon + \Delta t + \Delta x^2),
\end{aligned} \tag{7.4.7}$$

where we used $g^n = c + \mathcal{O}(\epsilon)$ for $n \ge 1$.

Finally, noticing that

$$\rho^{n+1} - \int M^n g^{n+1}\,\mathrm{d}\boldsymbol{v} = \int(M^{n+1}-M^n)g^{n+1}\,\mathrm{d}\boldsymbol{v} = c\int(M^{n+1}-M^n)\,\mathrm{d}\boldsymbol{v} + \mathcal{O}(\epsilon) = \mathcal{O}(\epsilon), \tag{7.4.8}$$

and using $J^n = \rho^n E^n + \mathcal{O}(\epsilon)$ for $n \geq 1$ (7.4.7) becomes

$$\frac{\rho^{n+1} - \rho^n}{\Delta t} = -\nabla_x \cdot (\rho^n E^n) + \mathcal{O}\left(\epsilon + \Delta t + \Delta x^2 + \frac{\epsilon}{\Delta t}\right). \tag{7.4.9}$$

(7.4.2) becomes

$$\frac{E^{n+1} - E^n}{\Delta t} = -\rho^n E^n + \mathcal{O}(\epsilon). \tag{7.4.10}$$

Equations (7.4.9) and (7.4.10) form a first-order in time discretization of (7.1.10) as $\epsilon \to 0$, as expected.

7.5 Numerical results

Example 7.5.1. To verify that the method described efficiently captures the asymptotic limit (7.1.10), we examine the behavior of the singular values of the low-rank solution for very small values of ϵ. The singular values are simply the diagonal entries of S, which we normalize by the largest singular value σ_1.

We consider two test cases in 1 dimension. In both test cases, ϵ is set to $\epsilon = 1 \times 10^{-6}$. The first is a solution beginning in local equilibrium, given by

$$f(x, v, 0) = \frac{\rho_0(x)}{\sqrt{2\pi}} e^{-\frac{|v-E_0|^2}{2}}, \tag{7.5.1}$$

$$\rho_0(x) = \frac{\sqrt{2\pi}}{2}(2 + \cos(2\pi x)), \tag{7.5.2}$$

where E_0 satisfies (7.1.3) with the ion background density given by

$$\eta(x) = \frac{\sqrt{2\pi}}{1.2661} e^{\cos(2\pi x)}.$$

We evolve the initial condition with a fixed rank of $r = 5$ until time $t = 0.01$, long enough to demonstrate that the asymptotic limit is captured. $N_x = N_v = 128$ grid points are used in each direction, and the timestep is chosen as $\Delta t = 3.9 \times 10^{-4}$. The evolution of the singular values is shown in Figure 7.1a.

The second test initializes f as a pair of counterstreaming beams:

$$f(x, v, 0) = \frac{\rho_0(x)}{2\sqrt{2\pi}} \left[e^{-\frac{|v-1.5|^2}{2}} + e^{-\frac{|v+1.5|^2}{2}} \right], \tag{7.5.3}$$

$$\rho_0(x) = \sqrt{2\pi}(2 + \cos(2\pi x)) \tag{7.5.4}$$

This is evolved with the same parameters as described above. Equation (7.5.3) is a highly non-equilibrium initial condition. Nevertheless, the solution quickly arrives at a numerical rank of 1, as demonstrated by Figure 7.1b.

Example 7.5.2. The efficiency benefits of the dynamical low-rank algorithm are most keenly felt in two

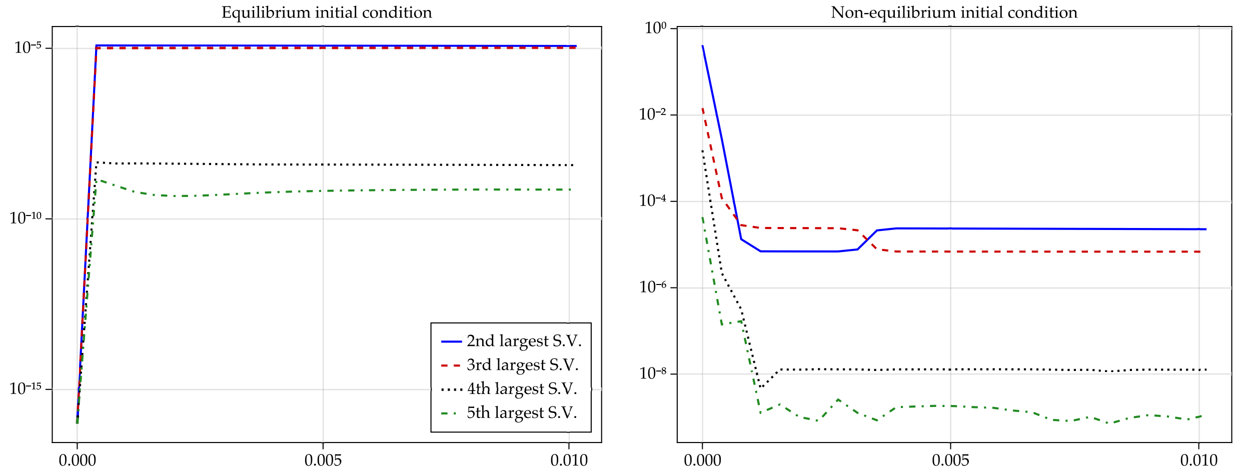

(a) Initial condition given by (7.5.1).

(b) Initial condition given by (7.5.3).

Figure 7.1: Singular value history of the simulations described in Example 7.5.1. Plots illustrate the ratio of singular values σ_2 through σ_5 normalized by σ_1. In both cases, σ_2/σ_1 quickly arrives at roughly 1×10^{-5}, indicating that the solution is of numerical rank 1.

dimensions and higher. To illustrate the algorithm's efficiency and demonstrate that it can capture both fluid and kinetic phenomena, we perform a two-dimensional (2D2V) test. The test case involves a “pulse” of electron density climbing an electrostatic potential hill. We initialize the background density η uniformly everywhere in the domain, except for a band through the center of the domain where it is set to zero. This creates a region of negative charge density through which the electron fluid cannot pass unless it has enough inertia to do so. In the fluid limit, inertial forces are negligible and the electrons are confined by the potential, whereas in the kinetic regime the fluid has sufficient momentum to pass over the barrier.

The spatial domain is a doubly periodic rectangular domain,

$$\Omega = [0,1]^2.$$

The background density η is initialized piecewise constant,

$$\eta(x,y) = \begin{cases} \eta_0 & x \le 0.55, x > 0.7 \\ 0 & 0.55 < x \le 0.7 \end{cases},$$

where

$$\eta_0 = \int_\Omega \rho_0 \, \mathrm{d}\boldsymbol{x}.$$

We use an elongated Gaussian initial density centered to the left of the potential hill and oriented obliquely

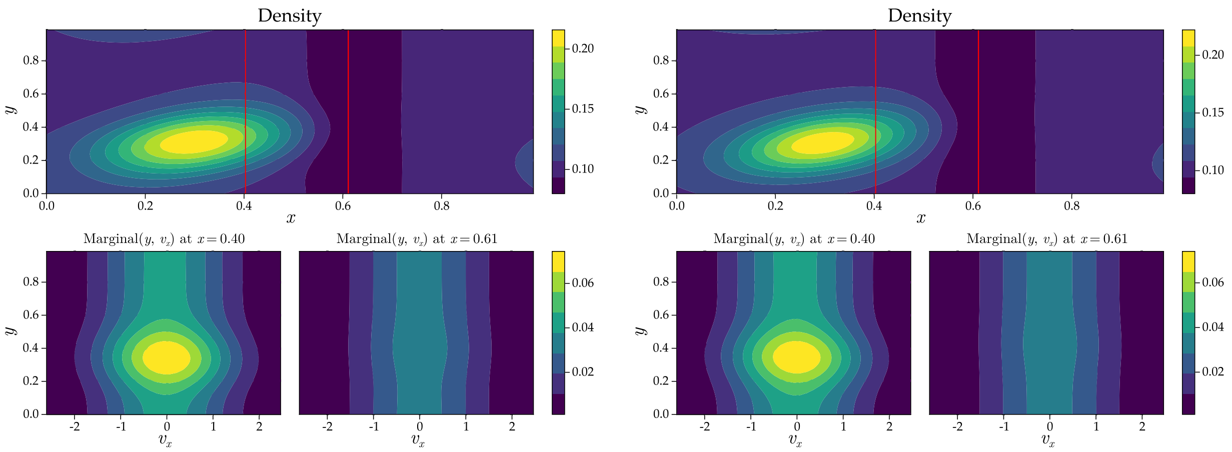

(a) Low rank solver, $\epsilon = 0.01$, $r = 10$.

(b) Full tensor solver, $\epsilon = 0.01$

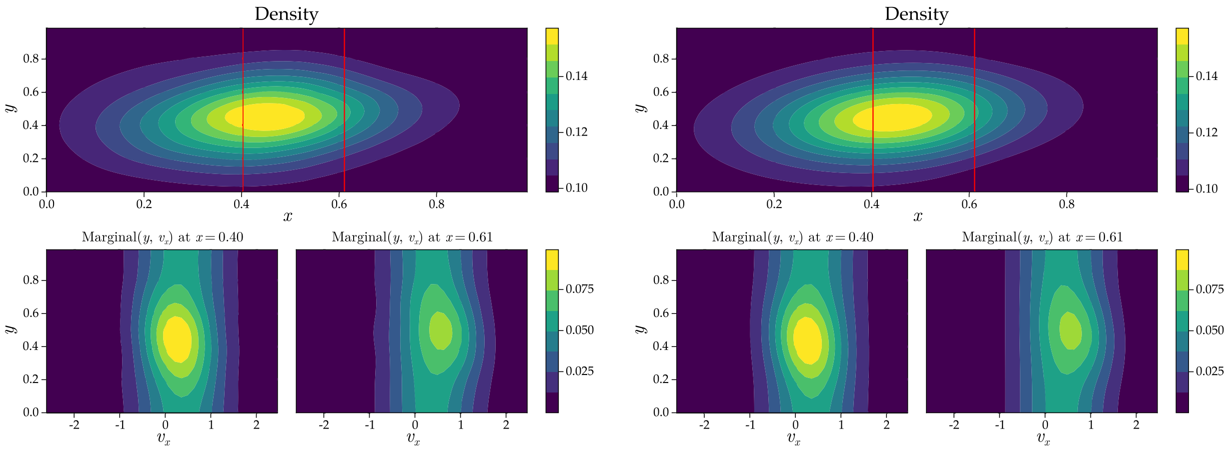

(c) Low rank solver, $\epsilon = 1.0$, $r = 15$.

(d) Full tensor solver, $\epsilon = 1.0$

Figure 7.2: Results of the electrostatic potential hill problem at $T = 0.35$, showing low-rank (left) and full tensor (right) solutions. The red lines indicate x coordinates of the marginal distributions plotted for each case. The results show good agreement between the low-rank and full tensor solutions for both the kinetic and fluid problems, indicating that our low-rank solver can handle both regimes.

to the grid:

$$\rho_0(\boldsymbol{x}, t) = 0.1 + \frac{3 \times 10^{-4}}{2\pi|\boldsymbol{\Sigma}|} e^{-\frac{(\boldsymbol{x}-\boldsymbol{x}_0)^T \boldsymbol{\Sigma}^{-1} (\boldsymbol{x}-\boldsymbol{x}_0)}{2}}, \quad \boldsymbol{x} \in \Omega.$$

Here, $\boldsymbol{x}_0 = [0.3, 0.3]$ and $\boldsymbol{\Sigma} = \mathbf{R}\Lambda\mathbf{R}^{-1}$, Λ is a diagonal matrix with entries $[0.006, 0.03]$, and $\mathbf{R}$ is a rotation matrix through an angle of $\pi/4$.

The two regimes, fluid and kinetic, are illustrated with the values $\epsilon = 0.01$ and $\epsilon = 1.0$ respectively. The low-rank solver uses a rank of $r = 10$ for the fluid case, and $r = 15$ for the kinetic case. The computational domain uses $N = 72$ grid points for each of x, y, v_x, v_y. We use a fixed time step of $\Delta t = 6.9 \times 10^{-4}$. Figure 7.2 plots the results at time $T = 0.35$.

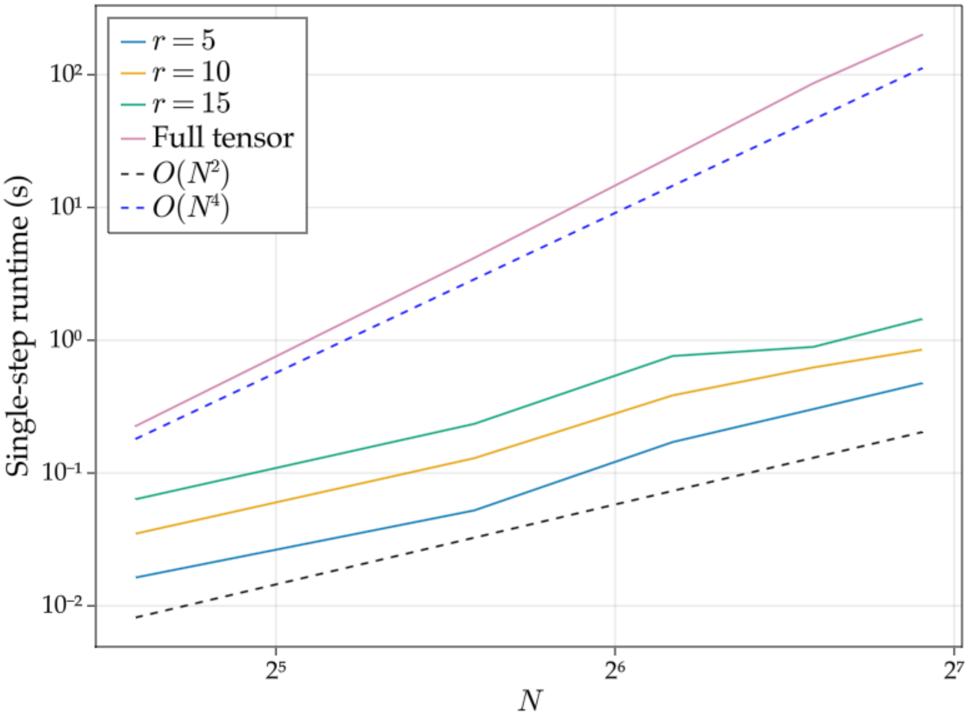

Figure 7.3: Runtime of a single timestep of the full tensor solver and the low-rank solver with ranks $r = \{5, 10, 15\}$ on the potential hill problem of Example 7.5.2. The spatial domain is discretized with N grid points in each dimension. The low-rank solver reduces computational complexity from $\mathcal{O}(N^4)$ to $\mathcal{O}(N^2)$.

7.6 Discussion of Forward Euler S step

In this section we motivate the choice of a Forward Euler timestep for the stiff part of the S flow in (7.2.34). Recall that the S step evolves just the singular values of the solution backwards in time, and comes in between the K and L steps:

$$\begin{aligned} K \text{ step:} \quad & g^n \xrightarrow{X^n \to X^{n+1},\ S^n \to S^1} g^1 \\ S \text{ step:} \quad & g^1 \xrightarrow{S^1 \to S^2} g^2 \\ L \text{ step:} \quad & g^2 \xrightarrow{V^n \to V^{n+1},\ S^2 \to S^{n+1}} g^{n+1} \end{aligned} \tag{7.6.1}$$

Because the S step is a diffusion equation solved backwards in time, strictly speaking it is an ill-posed ODE. However, empirically we find that an IMEX step for K followed by a Forwards Euler step for S is stable. To motivate this with a heuristic argument, consider an initial condition that is uniform in space, with vanishing electric field and current. We also assume that the solution begins in local equilibrium. In terms of our low-rank method, we take

$$g(x, v, t_0) = 1, \quad E(x) = J(x) = 0. \tag{7.6.2}$$

The low-rank decomposition of g gives us $S_{ij} = \delta_{i1}\delta_{j1}$, i.e. S_{11} is the only nonzero entry of S. With no spatial dependence, all but the collisional term of (7.2.26) drop out, and we are left with

$$\partial_t K_j = \frac{1}{\epsilon} \sum_l d^1_{jl} K_l, \tag{7.6.3}$$

where d^1_{jl} is defined in (7.2.21). The time evolution equation (7.2.33) for S also simplifies:

$$\partial_t S_{ij} = -\frac{1}{\epsilon} \sum_{kl} \delta_{ik} d^1_{jl} S_{kl} = -\frac{1}{\epsilon} \sum_l d^1_{jl} S_{il}. \tag{7.6.4}$$

It is useful to rewrite these equations in matrix form. Define the matrices $\mathbf{X} \in \mathbb{R}^{N_x \times r}$, $\mathbf{S} \in \mathbb{R}^{r \times r}$, $\mathbf{V} \in \mathbb{R}^{N_v \times r}$. Then $g = \mathbf{X}\mathbf{S}\mathbf{V}^T = \mathbf{K}\mathbf{V}^T$. Further define $\mathbf{D} = \{d^1_{jl}\}$. The K flow and S flow are given by

$$\partial_t \mathbf{K} = \frac{1}{\epsilon} \mathbf{K}\mathbf{D}^T, \tag{7.6.5}$$

$$\partial_t \mathbf{S} = -\frac{1}{\epsilon} \mathbf{S}\mathbf{D}^T. \tag{7.6.6}$$

During the K step and the S step, the respective time derivatives of g are equal and opposite:

$$\text{K step:} \qquad \partial_t g = (\partial_t \mathbf{K})\mathbf{V}^T = \frac{1}{\epsilon}\mathbf{K}\mathbf{D}^T\mathbf{V}^T = \frac{1}{\epsilon}\boldsymbol{X}\mathbf{S}\mathbf{D}^T\mathbf{V}^T, \tag{7.6.7}$$

$$\text{S step:} \qquad \partial_t g = \mathbf{X}(\partial_t \mathbf{S})\mathbf{V}^T = -\frac{1}{\epsilon}\mathbf{X}\mathbf{S}\mathbf{D}^T\mathbf{V}^T. \tag{7.6.8}$$

It follows that at the continuous level, our low-rank approximation has the property that $g^n = g^2$ for spatially homogeneous starting point g^n. We choose our time discretization to preserve this invariant. Using the backward Euler for the K flow results in

$$\mathbf{K}^{n+1} = \mathbf{K}^n \left(I - \frac{\Delta t}{\epsilon}\mathbf{D}^T\right)^{-1}, \tag{7.6.9}$$

while using the forward Euler for the S flow results in

$$\mathbf{S}^2 = \mathbf{S}^1 \left(I - \frac{\Delta t}{\epsilon}\mathbf{D}^T\right). \tag{7.6.10}$$

Therefore,

$$\begin{aligned}
g^2 &= \mathbf{X}^{n+1}\mathbf{S}^2(\mathbf{V}^n)^T \\
&= \mathbf{X}^{n+1}\mathbf{S}^1\left(I - \frac{\Delta t}{\epsilon}\mathbf{D}^T\right)(\mathbf{V}^n)^T \\
&= \mathbf{K}^{n+1}\left(I - \frac{\Delta t}{\epsilon}\mathbf{D}^T\right)(\mathbf{V}^n)^T \\
&= \mathbf{K}^{n}\left(I - \frac{\Delta t}{\epsilon}\mathbf{D}^T\right)^{-1}\left(I - \frac{\Delta t}{\epsilon}\mathbf{D}^T\right)(\mathbf{V}^n)^T \\
&= g^n.
\end{aligned}$$

While we have exact cancellation of the K and S flows for this spatially homogeneous equilibrium, each of these substeps is quite large when taken individually. The stiffness of the ϵ^{-1} term means that it is quite important to preserve this cancellation at the numerical level, otherwise the method is unable to hold even a spatially homogeneous equilibrium. For example, if one uses an IMEX step for the S flow, one finds

$$\begin{aligned}
g^2 &= \mathbf{K}^{n}\left(I - \frac{\Delta t}{\epsilon}\mathbf{D}^T\right)^{-1}\left(I + \frac{\Delta t}{\epsilon}\mathbf{D}^T\right)^{-1}(\mathbf{V}^n)^T \\
&= \mathbf{K}^{n}\left(I - \left(\frac{\Delta t}{\epsilon}\right)^2(\mathbf{D}^T)^2\right)^{-1}(\mathbf{V}^n)^T \qquad\qquad \neq g^n.
\end{aligned}$$

The preceding argument is not a rigorous justification of our choice of timestepping scheme for general solutions g. We simply wish to highlight one subtle numerical aspect of the projector-splitting approach which implementors should be aware of.

7.6.1 Comparison to the BGK operator

In [33], the authors successfully used an IMEX step to advance the S flow. The issues raised above do not arise for the BGK-type operators considered there, as we demonstrate here with a simple example. A spatially homogeneous equation with BGK-type collision operator is

$$\partial_t f = \frac{1}{\epsilon}(M - f),$$

or, using the fact that M is constant for a spatially homogeneous problem,

$$\partial_t g = \frac{1}{\epsilon}(1 - g).$$

Projecting this onto the low-rank approximation gives the following subflows for K and S:

$$\partial_t K_j = \frac{1}{\epsilon}(\langle V_j \rangle_v - K_j), \quad \partial_t S_{ij} = -\frac{1}{\epsilon}(\langle X_i V_j \rangle_{xv} - S_{ij}).$$

An IMEX (backwards Euler) step for each of these subflows will give

$$\begin{aligned} \mathbf{K}^{n+1} &= \left(1 + \frac{\Delta t}{\epsilon}\right)^{-1} \left(\mathbf{K}^n + \frac{\Delta t}{\epsilon}\langle (\mathbf{V}^n)^T \rangle_v\right), \\ \mathbf{S}^2 &= \left(1 - \frac{\Delta t}{\epsilon}\right)^{-1} \left(\mathbf{S}^1 - \frac{\Delta t}{\epsilon}\langle \mathbf{X}^{n+1} (\mathbf{V}^n)^T \rangle_{xv}\right). \end{aligned}$$

Plugging these into the expression for g^2, we find

$$\begin{aligned} g^2 &= \mathbf{X}^{n+1}\mathbf{S}^2(\mathbf{V}^n)^T \\ &= \mathbf{X}^{n+1}\left(\mathbf{S}^1 - \frac{\Delta t}{\epsilon}\langle \mathbf{X}^{n+1}(\mathbf{V}^n)^T\rangle_{xv}\right)\left(1 - \frac{\Delta t}{\epsilon}\right)^{-1}(\mathbf{V}^n)^T \\ &= \mathbf{K}^{n+1}\left(1 - \frac{\Delta t}{\epsilon}\right)^{-1}(\mathbf{V}^n)^T - \mathbf{X}^{n+1}\left(\frac{\epsilon}{\Delta t} - 1\right)^{-1}\langle \mathbf{X}^{n+1}(\mathbf{V}^n)^T\rangle_{xv}(\mathbf{V}^n)^T \\ &= \mathbf{X}^{n+1}\langle 1, \mathbf{X}^{n+1}(\mathbf{V}^n)^T\rangle_{xv}(\mathbf{V}^n)^T + O\left(\frac{\epsilon}{\Delta t}\right). \end{aligned}$$

Note that we have dropped terms of order $\epsilon/\Delta t$, to illustrate that g is driven to within $\epsilon/\Delta t$ of its equilibrium value, which is 1 (projected onto the low-rank bases). Because the BGK operator on g is affine rather than linear there is no cancellation, but the IMEX approach for both flows poses no problems in the $\epsilon \to 0$ limit.

Chapter 8

A conservative DLR algorithm for the Vlasov-Lenard-Bernstein equation

The dynamical low-rank algorithms described in Example 6.3.1 and Chapter 7 illustrate the efficiency gains that are possible through the DLR framework. The number of degrees of freedom in a DLR discretization of the Vlasov equation can be orders of magnitude lower than a similar full-tensor discretization of f. The computational efficiency gains are proportional.

However, DLR methods do not necessarily preserve important physical structure of the kinetic equation. Chapter 7 showed how we can modify the DLR method for a Vlasov-type equation to capture the strongly collisional limit with only a small number of ranks, thereby preserving the local Maxwellian structure of the fluid asymptotic limit. In this chapter we are concerned with another important piece of physical structure, which is the conservation of mass, momentum and energy.

As demonstrated by the derivation of the five-moment fluid model, the Vlasov equation conserves the mass, momentum and energy of the plasma [1]. Respecting these conservation laws is an important goal for a usable numerical discretization of the Vlasov equation. However, a standard DLR approximation such as Example 6.3.1 will not be conservative. The truncation performed by the low-rank approximation does not respect the desired conservation properties: the truncated ranks may carry non-negligible amounts of mass, momentum or energy.

As such, there is considerable interest in deriving modified DLR schemes that conserve mass, momentum and energy of f. A "quasi-conservative" method based on Lagrange multipliers was introduced in [37]. This approach only achieves global conservation however, implying the possibility of unphysically fast numerical

---

[1] The Vlasov equation actually conserves an infinite number of moments, but a collisional Vlasov-Boltzmann or Vlasov-Landau equation conserves only these three.

transport of mass and energy through the solution. Recognizing this shortcoming, locally conservative DLR methods were quickly developed by the community. A first-principles method of achieving conservation was laid out in [34, 38]. These derivations force the velocity basis to contain the functions $(1, v, |v|^2/2)$ by modifying the DLR Galerkin condition (6.1.9). The second of these papers, [38], showed how to achieve this in the context of the Basis-update & Galerkin (BUG) integrator, thereby obtaining a conservative DLR algorithm that is robust to the presence of small singular values. More recently, it was observed that the rank-adaptive BUG integrator is conservative if equipped with a conservative rank-truncation algorithm [35].

In [84], the DLR method was used to evolve the high-order component of a high-order/low-order (HOLO) scheme for the radiative transport equation. This approach employs the DLR method as a moment closure for a low-order fluid system. A least-squares projection is applied to ensure that the DLR solution's conserved moments remain close to the fluid system's solution. The resulting combination can be viewed as a conservative DLR method. The step-truncation family of methods, which project the solution rather than its time derivative, are also amenable to conservative modifications. For example, conservative projection methods have been used to attain locally conservative low-rank solutions to the Vlasov equation [47, 46, 48].

The method we discuss here is perhaps closest to work done in [73]. There, the authors show how to use a spectral Legendre discretization of a kinetic extension of the shallow-water equations to combine conservation properties with the BUG integrator. A low-rank *ansatz* is applied to the trailing "microscopic" modes while evolving the leading modes using a standard conservative discretization. However, as formulated, the method is limited to a modal discretization of the phase space coordinate.

In this chapter, we present a new way of obtaining a conservative DLR method for kinetic equations, focusing on the Vlasov-Lenard-Bernstein equation. The work described here has been published in [20]. Our method is based on a novel macro-micro decomposition that operates at the equation level, and is designed to be amenable to low-rank approximation. As in [73], the macroscopic part of the solution may be evolved using standard conservative discretizations, while the microscopic part is evolved using a DLR integrator.

There are several benefits to our formulation. By applying the macro-micro decomposition at the equation level rather than following a modal discretization, we are able to apply any desired velocity space discretization. This has particular benefits for the plasma applications considered here. Shock-capturing discretizations in velocity space such as Discontinuous Galerkin are very popular for their ability to resolve fine phase space structures [59, 51]. Moreover, our method is compatible with any of the variety of DLR integrators, including BUG and the projector-splitting integrator (PSI). Here, we apply the PSI to achieve formal second-order accuracy in time.

The rest of this chapter is organized as follows. In the following section we re-introduce the Vlasov equation with Lenard-Bernstein collision operator, and formally state the local conservation laws we aim to replicate in our discretization. In Section 8.2 we describe our novel macro-micro decomposition and derive the equations of evolution for each part. Section 8.3 presents a first- and second-order time integrator for the macro-micro decomposition and proves exact local conservation. In Section 8.4 we briefly describe a pair of spatial discretizations, and in Section 8.6 we present numerical results on a handful of benchmark problems.

8.1 Conservation properties of the Vlasov equation with Lenard-Bernstein collisions

In this chapter we restrict our attention to a single-species Vlasov equation with the Lenard-Bernstein collision operator. Taking the trivial unit normalization,

$$\omega_p \tau = \omega_c \tau = \nu_p \tau = Z_a = A_a = 1,$$

equation (2.4.2) becomes

$$\partial_t f + \boldsymbol{v} \cdot \nabla_{\boldsymbol{x}} f + \boldsymbol{E} \cdot \nabla_{\boldsymbol{v}} f = C(f), \tag{8.1.1}$$

where $t > 0$, $\boldsymbol{x} \in \Omega_x \subset \mathbb{R}^d$, and $\boldsymbol{v} \in \mathbb{R}^d$. The physical interpretation of a single-species Vlasov equation such as (8.1.1) is that it describes the motion of the electrons against a static ion background. In this case, we note that our charge convention is reversed, so that the electrons are given a positive unit charge.

We couple (8.1.1) with either Gauss's law,

$$\nabla \cdot \boldsymbol{E}(\boldsymbol{x}, t) = \rho(\boldsymbol{x}, t) - \rho_0, \quad \rho(\boldsymbol{x}, t) = \int_{\mathbb{R}^d} f \, \mathrm{d}\boldsymbol{v},$$

where ρ_0 is a uniform background density given by $\int_{\Omega_x} \rho(\boldsymbol{x}, t) - \rho_0 \, \mathrm{d}\boldsymbol{x} = 0$, or by Ampère's law:

$$\partial_t \boldsymbol{E}(\boldsymbol{x}, t) = -\boldsymbol{J}(\boldsymbol{x}, t), \quad \boldsymbol{J}(\boldsymbol{x}, t) = \int_{\mathbb{R}^d} \boldsymbol{v} f \, \mathrm{d}\boldsymbol{v}. \tag{8.1.2}$$

The Lenard-Bernstein collision operator for a single species is given by

$$C(f) = \nu \nabla_{\boldsymbol{v}} \cdot (T \nabla_{\boldsymbol{v}} f + (\boldsymbol{v} - \boldsymbol{u}) f),$$

where the temperature T and drift velocity $\boldsymbol{u}$ are defined in the usual way:

$$\boldsymbol{u}(\boldsymbol{x}, t) = \frac{1}{\rho} \int_{\mathbb{R}^d} \boldsymbol{v} f \, \mathrm{d}\boldsymbol{v}, \quad T(\boldsymbol{x}, t) = \frac{1}{d\rho} \int_{\mathbb{R}^d} |\boldsymbol{v} - \boldsymbol{u}|^2 f \, \mathrm{d}\boldsymbol{v}.$$

The collision operator C satisfies the trio of conservation identities

$$\int_{\mathbb{R}^d} \boldsymbol{\phi}(\boldsymbol{v}) C(f)\,\mathrm{d}\boldsymbol{v} = \mathbf{0},$$

where $\boldsymbol{\phi}(\boldsymbol{v}) = (1, \boldsymbol{v}, |\boldsymbol{v}|^2/2)^T$ is the vector of collision invariants. These collision invariants correspond to the conservation of mass, momentum and energy in interparticle collisions.

Each element of $\boldsymbol{\phi}$ admits a local conservation law. Define the momentum density $\boldsymbol{J}$, the kinetic energy density κ, and the total energy density e as

$$\rho\boldsymbol{u}(\boldsymbol{x}, t) = \int_{\mathbb{R}^d} \boldsymbol{v} f\,\mathrm{d}\boldsymbol{v}, \tag{8.1.3}$$

$$\kappa(\boldsymbol{x}, t) = \int_{\mathbb{R}^d} \frac{|\boldsymbol{v}|^2}{2} f\,\mathrm{d}\boldsymbol{v}, \tag{8.1.4}$$

$$e(\boldsymbol{x}, t) = \kappa + \frac{|\boldsymbol{E}|^2}{2}. \tag{8.1.5}$$

Note that we have denoted the momentum by $\rho\boldsymbol{u}$, while previously we denoted the same quantity on the right-hand side of Ampère's law (8.1.2) by $\boldsymbol{J}$. In the context of the single-species Vlasov equation with unit normalization, (8.1.1), the current and momentum densities are equal, although this is not true in general.

Taking the moments of (8.1.1) weighted by each component of $\boldsymbol{\phi}(\boldsymbol{v})$ gives a system of three local conservation laws for $(\rho, \rho\boldsymbol{u}, e)^T$:

$$\partial_t \rho + \nabla \cdot (\rho\boldsymbol{u}) = 0, \tag{8.1.6}$$

$$\partial_t (\rho\boldsymbol{u}) + \nabla \cdot \sigma = \rho\boldsymbol{E}, \tag{8.1.7}$$

$$\partial_t e + \nabla \mathbf{q} = 0, \tag{8.1.8}$$

where $\sigma = \int_{\mathbb{R}^d} (\boldsymbol{v} \otimes \boldsymbol{v}) f\,\mathrm{d}\boldsymbol{v}$ and $\mathbf{q} = \frac{1}{2}\int_{\mathbb{R}^d} \boldsymbol{v}|\boldsymbol{v}|^2 f\,\mathrm{d}\boldsymbol{v}$.

8.2 A macro-micro decomposition of the Vlasov equation

By macro-micro decompositions, we refer to a family of methods that use a decomposition of the form

$$f(\boldsymbol{x}, \boldsymbol{v}, t) = \mathcal{N}(\boldsymbol{x}, \boldsymbol{v}, t) + g(\boldsymbol{x}, \boldsymbol{v}, t),$$

where $\mathcal{N}$ is chosen to share its first $d + 2$ moments with f, while the same moments of g vanish. Given such a decomposition, the strategy is to evolve $\mathcal{N}$ as accurately as possible using standard conservative discretization methods. Meanwhile, g does not contribute to the mass, momentum or energy density of the solution, so we may evolve it at a precision dictated by the kinetic physics of the problem. The macro-micro decomposition splits a high-dimensional problem into two parts: a lower-dimensional problem which can be solved for $\mathcal{N}$ conservatively, and a high-dimensional problem for g, to which we can apply a coarser discretization or more sophisticated dimension reduction techniques.

The most natural and widely-known macro-micro decomposition for collisional kinetic equations is given in [4]. The idea of this decomposition is to choose a Maxwellian $\mathcal{M}$ as the "macro" component:

$$\mathcal{N} = \mathcal{M}(\boldsymbol{x}, \boldsymbol{v}, t) = \frac{\rho(\boldsymbol{x}, t)}{(2\pi T(\boldsymbol{x}, t))^{d/2}} e^{-\frac{|\boldsymbol{v}-\boldsymbol{u}(\boldsymbol{x},t)|^2}{2T(\boldsymbol{x},t)}}.$$

For strongly collisional problems, such a decomposition can anticipate that the remainder $g = f - \mathcal{M}$ will be small compared to $\mathcal{M}$. However, this decomposition is much less appealing when collisions are weak. Moreover, it is not favored by the dynamical low-rank method. A key step in a DLR scheme is the projection onto a low-rank approximate solution manifold, such as (6.3.11). When a Maxwellian-based macro-micro decomposition is used, the projection step requires computing an integral such as

$$\left(\frac{\rho(\boldsymbol{x}, t)}{2\pi T(\boldsymbol{x}, t)}\right)^{d/2} \int_{\mathbb{R}^d} V_j(\boldsymbol{v}) e^{-\frac{|\boldsymbol{v}-\boldsymbol{u}(\boldsymbol{x},t)|^2}{2T(\boldsymbol{x},t)}} \,\mathrm{d}\boldsymbol{v}. \tag{8.2.1}$$

In Chapter 7, we computed a similar integral for an isothermal Maxwellian using a convolution trick. For non-isothermal flows, the integral (8.2.1) cannot be computed efficiently.

Setting aside the numerical difficulties, a Maxwellian-based macro-micro decomposition does not hold much appeal for us. For plasma applications near the electron scale, collisions are typically much weaker than neutral gases and phase space can exhibit highly non-equilibrium features. Therefore, we do not expect a decomposition such as [4], nor a Maxwellian-centered DLR scheme such as [33, 19] to be advantageous. Rather, to avoid the difficulties presented by integrals such as (8.2.1), we derive a macro-micro decomposition whose macro component $\mathcal{N}$ has a fixed rank of $d + 2$.

8.2.1 A novel macro-micro decomposition in 1D1V

The remainder of this chapter considers a fixed dimension $d = 1$. That is, we consider the "1D1V" case. Our notation will use unbolded letters for the phase space coordinates x, v and the electric field E. As this point we also allow a truncated velocity space: $v \in \Omega_v \subseteq \mathbb{R}$. The one-dimensional, truncated Vlasov equation is therefore

$$\partial_t f + v\partial_x f + E\partial_v f = C(f), \quad t > 0, \quad x \in \Omega_x \subset \mathbb{R}, \quad v \in \Omega_v \subseteq \mathbb{R}. \tag{8.2.2}$$

Denote the phase space domain by $\Omega = \Omega_x \times \Omega_v$. For simplicity, we impose periodic boundary conditions in x. No boundary conditions are required for an unbounded velocity domain, although we must assume that f decays sufficiently quickly as $|v| \to \infty$. For a bounded velocity domain, we make the same rapid decay assumption, such that f and its derivatives are negligible at the velocity boundary.

For completeness, we redefine the fluid variables in terms of integrals over Ω_v:

$$\rho = \langle f \rangle_v , \quad J = \rho u = \langle v f \rangle_v , \quad \kappa = \left\langle \frac{|v|^2}{2} f \right\rangle_v , \quad e = \kappa + \frac{|E|^2}{2}, \tag{8.2.3}$$

$$u = \frac{\rho u}{\rho}, \quad T = \frac{1}{\rho} \left\langle |v-u|^2 f \right\rangle_v , \tag{8.2.4}$$

where $\langle \cdot \rangle_v = \int_{\Omega_v} \cdot \, \mathrm{d}v$. Using the fast decay of f and its derivatives, we can again show

$$\langle \boldsymbol{\phi}(v) C(f) \rangle_v = \mathbf{0}.$$

Therefore, the local conservation laws (8.1.6)-(8.1.8) still hold for (8.2.2).

Our macro-micro decomposition is based on orthogonal projection in an inner product space over Ω_v. To construct such an inner product space over a possibly unbounded domain, we will use a weight function which we denote $w(v) : \Omega_v \mapsto \mathbb{R}$. The function $w(v)$ induces a pair of weighted inner products on Ω_v and Ω:

$$\langle g, h \rangle_{w^{-1}(v)} = \int_{\Omega_v} w^{-1}(v) g(v) h(v) \, \mathrm{d}v, \quad \langle g, h \rangle_{x, w^{-1}(v)} = \int_{\Omega} w^{-1}(v) g(x,v) h(x,v) \, \mathrm{d}x \, \mathrm{d}v.$$

The corresponding inner product spaces will be denoted by $L^2(\Omega_v, w^{-1})$ and $L^2(\Omega, w^{-1})$ respectively. From standard theory [43], the inner product $\langle \cdot, \cdot \rangle_{w^{-1}(v)}$ has an associated family of orthonormal polynomials, which we denote $p_n(v)$, satisfying the orthogonality relation

$$\langle w(v) p_n(v), w(v) p_m(v) \rangle_{w^{-1}(v)} = \int_{\Omega_v} w(v) p_n(v) p_m(v) \, \mathrm{d}v = \delta_{nm}, \tag{8.2.5}$$

along with a three-term recurrence relation

$$v p_n(v) = a_n p_{n+1}(v) + b_n p_n(v) + a_{n-1} p_{n-1}(v). \tag{8.2.6}$$

Because the polynomials are ordered by degree, we can also define three coefficients, d_{10}, d_{20}, d_{21} relating the derivatives of p_1 and p_2 by

$$p_1'(v) = d_{10} p_0(v), \quad p_2'(v) = d_{20} p_0(v) + d_{21} p_1(v). \tag{8.2.7}$$

The first three orthonormal polynomials, which we denote by $\mathbf{p}(v) = (p_0(v), p_1(v), p_2(v))^T$, are related to the collision invariants $\boldsymbol{\phi}(v)$ by a lower-triangular, invertible matrix Q:

$$\boldsymbol{\phi}(v) = \begin{pmatrix} 1 \\ v \\ |v|^2/2 \end{pmatrix} = \begin{pmatrix} q_{00} & & \\ q_{10} & q_{11} & \\ q_{20} & q_{21} & q_{22} \end{pmatrix} \begin{pmatrix} p_0(v) \\ p_1(v) \\ p_2(v) \end{pmatrix} = Q \begin{pmatrix} p_0(v) \\ p_1(v) \\ p_2(v) \end{pmatrix} = Q\mathbf{p}(v). \tag{8.2.8}$$

The span of $w(v)\boldsymbol{\phi}(v)$ plays an important role in our construction. We denote it by Φ:

$$\Phi = \operatorname{span}\left\{w(v), vw(v), |v|^2 w(v)/2\right\} = \operatorname{span}\{w(v)\mathbf{p}(v)\}. \tag{8.2.9}$$

Equation (8.2.9) also states that $w(v)\mathbf{p}(v)$ is a basis for Φ, which is assured by the existence of the matrix Q. In fact it is an orthogonal basis with respect to the inner product $\langle \cdot, \cdot \rangle_{w^{-1}(v)}$ since its elements are orthogonal per (8.2.5). Thus, the orthogonal projection onto Φ has the explicit formula

$$P_\Phi f = w(v)\mathbf{p}(v)^T \langle w(v)\mathbf{p}(v), f\rangle_{w^{-1}(v)} = w(v)\mathbf{p}(v)^T \langle \mathbf{p}(v) f\rangle_v, \tag{8.2.10}$$

where we have introduced notation for the unweighted integral $\langle \cdot \rangle_v = \int_{\Omega_v} \cdot \, \mathrm{d}v$. The orthogonal complement of P_Φ is $P_\Phi^\perp = I - P_\Phi$.

Having defined the subspace Φ and the associated projections, we can now state our proposed macro-micro decomposition: it is precisely the decomposition induced by P_Φ and its orthogonal complement,

$$f(x,v,t) = \underbrace{P_\Phi f(x,v,t)}_{\mathcal{N}} + \underbrace{P_\Phi^\perp f(x,v,t)}_{g}. \tag{8.2.11}$$

The function $\mathcal{N}$ has an explicit formula in terms of p_0, p_1, p_2 and their moments:

$$\mathcal{N}(x,v,t) = w(v)\mathbf{p}(v)^T \langle \mathbf{p}(v) f\rangle_v = w(v)[p_0(v)f_0(x,t) + p_1(v)f_1(x,t) + p_2(v)f_2(x,t)], \tag{8.2.12}$$

where f_n denotes the moment of f with respect to $p_n(v)$:

$$f_n(x,t) = \langle p_n(v) f(x,v,t)\rangle_v .$$

Using the orthogonality relation (8.2.5), it is easy to show that $\langle \mathbf{p}\mathcal{N}\rangle_v = \langle \mathbf{p} f\rangle_v$. Because $\mathbf{p}(v)$ and $\boldsymbol{\phi}(v)$ are related by the matrix Q, it also follows that $\mathcal{N}$ and f have the same moments with respect to $\boldsymbol{\phi}(v)$:

$$\langle \boldsymbol{\phi}\mathcal{N}\rangle_v = Q\langle \mathbf{p}(v)\mathcal{N}\rangle_v = Q\langle \mathbf{p}(v) f\rangle_v = \langle \boldsymbol{\phi} f\rangle_v .$$

This implies that $\langle \boldsymbol{\phi} g\rangle_v = 0$. That is, the mass, momentum, and kinetic energy density of the microscopic part g of the distribution function all vanish.

We now derive equations for the macroscopic and microscopic components of our decomposition. From equation (8.2.12), we can see that for the macroscopic component, it suffices to derive equations for f_0, f_1, f_2. This is easily done by taking the moments of (8.2.2) with respect to $\mathbf{p}(v)$. The result, after making use of the recurrence and derivative relations (8.2.6) and (8.2.7), is

$$\partial_t \underbrace{\begin{pmatrix} f_0 \\ f_1 \\ f_2 \end{pmatrix}}_{\mathbf{f}} + \underbrace{\begin{pmatrix} b_0 & a_0 & \\ a_0 & b_1 & a_1 \\ & a_1 & b_2 \end{pmatrix}}_{V} \partial_x \begin{pmatrix} f_0 \\ f_1 \\ f_2 \end{pmatrix} + \partial_x \begin{pmatrix} 0 \\ 0 \\ a_2 f_3 \end{pmatrix} + E \underbrace{\begin{pmatrix} 0 & & \\ -d_{10} & 0 & \\ -d_{20} & -d_{21} & 0 \end{pmatrix}}_{D_v} \begin{pmatrix} f_0 \\ f_1 \\ f_2 \end{pmatrix} = \mathbf{0}. \tag{8.2.13}$$

We have introduced the notation $\mathbf{f} = \langle \mathbf{p}(v) f \rangle_v$ for the first three moments of f, as well as V and D_v for the matrices appearing in (8.2.13). The right-hand side vanishes because the moments of $C(f)$ with respect to $\mathbf{p}(v)$ vanish. Note that (8.2.13) is an unclosed system of equations, due to the appearance of $\partial_x f_3$ in the third component. This is the manifestation of the usual moment closure problem for our scheme. The closure information must come from the microscopic component:

$$f_3 = \langle p_3(v) f \rangle_v = \langle p_3(v) g \rangle_v .$$

The initial-boundary value problem for $\mathbf{f}$ is therefore

$$\begin{cases} \partial_t \mathbf{f} + V \partial_x \mathbf{f} + a_2 \partial_x f_3 \mathbf{e}_2 + E(x,t) D_v \mathbf{f} = 0, & (x,t) \in \Omega_x \times (0,\infty), \\ \mathbf{f}(x,0) = \langle \mathbf{p}(v) f(x,v,0) \rangle_v , & x \in \Omega_x. \end{cases} \tag{8.2.14}$$

To obtain an expression for $\partial_t g$, we substitute $f = \mathcal{N} + g$ into (8.2.2), and apply the projection $P_\Phi^\perp$:

$$\begin{aligned} \partial_t g &= -\partial_t P_\Phi^\perp \mathcal{N} - P_\Phi^\perp (v \partial_x \mathcal{N} + E \partial_v \mathcal{N} + v \partial_x g + E \partial_v g) + P_\Phi^\perp C(\mathcal{N} + g) && (8.2.15) \\ &= -P_\Phi^\perp (v \partial_x \mathcal{N} + E \partial_v \mathcal{N} + v \partial_x g + E \partial_v g) + P_\Phi^\perp C(\mathcal{N} + g) && (8.2.16) \\ &\triangleq P_\Phi^\perp D[E, \mathcal{N}, g]. && (8.2.17) \end{aligned}$$

We must also modify the initial and boundary conditions to obtain an initial-boundary value problem for g:

$$\begin{cases} \partial_t g(x,v,t) = P_\Phi^\perp D[E, \mathcal{N}, g], & (x,v,t) \in \Omega \times (0,\infty), \\ g(x,v,0) = f(x,v,0) - \mathcal{N}(x,v,0), & (x,v) \in \Omega \\ g(x,v_b,t) = -\mathcal{N}(x,v_b,t), & (x,v,t) \in \Omega_x \times \partial\Omega_v \times (0,\infty). \end{cases} \tag{8.2.18}$$

Together, (8.2.14) and (8.2.18) constitute an exact macro-micro decomposition of (8.2.2). They are coupled by the appearance of f_3 in (8.2.14) on the one hand, and the appearance of $\mathcal{N}$ in (8.2.18) on the other.

8.3 Time discretization

We now discuss the discretization of (8.2.14) and (8.2.18) in time. The strategy is to use standard conservative techniques to discretize the macroscopic part $\mathcal{N}$, while the microscopic part g is evolved with a dynamical low-rank method. We begin by describing how to apply the projector-splitting integrator to (8.2.18). We then present a pair of integrators with different conservation properties:

- A first-order integrator which exactly conserves charge density and either momentum or energy density, depending on how f is coupled to E.
- A second-order integrator which exactly conserves charge and energy density by coupling to E via Ampère's law.

In each case we provide a proof of the claimed conservation property.

8.3.1 The projector-splitting integrator applied to g

A notable feature of our macro-micro decomposition is that it is independent of the details of the DLR scheme, and can in principle be combined with any appropriate DLR integrator. Here we choose the projector-splitting integrator (PSI), which can be formally extended to second-order via a Strang-splitting scheme. In prior work, locally conservative methods for the Vlasov equation have been described using the traditional [34] and BUG [38, 35] integrators. A globally conservative projector-splitting method based on Lagrange multipliers was described in [37]. To our knowledge this is the first locally conservative Vlasov scheme to use the PSI.

We introduce a slight modification of the PSI, which is that for projection along v, we use the weighted inner product $\langle \cdot \rangle_{w^{-1}(v)}$. The low-rank ansatz for g is

$$g(x, v, t) = \sum_{ij} X_i(x, t) S_{ij}(t) V_j(v, t), \tag{8.3.1}$$

where the basis functions X_i satisfy

$$X_i \in L^2(\Omega_x), \quad \langle X_i, X_k \rangle_x = \delta_{ik}.$$

The velocity basis functions V_j satisfy a similar condition for the weighted inner product:

$$V_j \in L^2(\Omega_v, w^{-1}(v)), \quad \langle V_j, V_l \rangle_{w^{-1}(v)} = \delta_{jl}.$$

As described in Section 6.2.1, the evolution of (8.3.1) can be obtained by projecting (8.2.15) onto the subspaces spanned by X_i and V_j, taking care to use the weighted inner product for projections in v. That is,

$$\begin{aligned}\partial_t g = &\sum_j \left\langle V_j, P_\Phi^\perp D[E,\mathcal{N},g]\right\rangle_{w^{-1}(v)} V_j - \sum_{ij} X_i \left\langle X_i V_j, P_\Phi^\perp D[E,\mathcal{N},g]\right\rangle_{x,w^{-1}(v)} V_j \\ &+ \sum_i X_i \left\langle X_i, P_\Phi^\perp D[E,\mathcal{N},g]\right\rangle_x .\end{aligned}$$

Following the standard notation, we define the auxiliary bases K_j and L_i via

$$K_j(x,t) = \sum_i X_i(x,t) S_{ij}(t), \quad L_i(v,t) = \sum_j S_{ij}(t) V_j(v,t). \tag{8.3.2}$$

In terms of the auxiliary bases,

$$\partial_t g = \sum_j \partial_t K_j V_j + \sum_{ij} X_i \partial_t S_{ij} V_j + \sum_i X_i \partial_t L_i, \tag{8.3.3}$$

with

$$\partial_t K_j = \left\langle V_j, P_\Phi^\perp D[E,\mathcal{N},g]\right\rangle_{w^{-1}(v)}, \tag{8.3.4}$$

$$\partial_t S_{ij} = -\left\langle X_i V_j, P_\Phi^\perp D[E,\mathcal{N},g]\right\rangle_{x,w^{-1}(v)}, \tag{8.3.5}$$

$$\partial_t L_i = \left\langle X_i, P_\Phi^\perp D[E,\mathcal{N},g]\right\rangle_x . \tag{8.3.6}$$

8.3.2 First-order integrator

In this section we describe a first-order time integrator which exactly conserves charge, and either momentum or energy density. The algorithm computes the following time advance between times t^n and $t^{n+1} = t^n + \Delta t$:

$$\begin{pmatrix} E^n \\ f_0^n \\ f_1^n \\ f_2^n \\ X^n \\ S^n \\ V^n \end{pmatrix} \mapsto \begin{pmatrix} E^{n+1} \\ f_0^{n+1} \\ f_1^{n+1} \\ f_2^{n+1} \\ X^{n+1} \\ S^{n+1} \\ V^{n+1} \end{pmatrix}.$$

1. **Calculate the electric field**: depending on the choice of conservation property, follow one of these

two procedures:

(a) **For momentum conservation**: Determine E^* from Gauss's law, by solving the following Poisson equation for the electric potential φ^n:

$$\partial_x^2 \varphi^n = -(\rho^n - \rho_0), \quad \rho^n = q_{00} f_0^n. \tag{8.3.7}$$

Then the electric field E^* to be used is

$$E^* = E^n = -\partial_x \varphi^n \tag{8.3.8}$$

(b) **For energy conservation**: Perform a Forward Euler step of Ampère's law (8.1.2) to obtain E^{n+1}:

$$\frac{E^{n+1} - E^n}{\Delta t} = -J^n, \tag{8.3.9}$$

where $J^n = \int_{\Omega_v} v f^n \,\mathrm{d}v = q_{10} f_0^n + q_{11} f_1^n$. The electric field to be used in the current timestep is the time-centered value

$$E^* = \frac{E^{n+1} - E^n}{2}.$$

2. **Advance conserved quantities**: Perform a single Forward Euler step of (8.2.13):

$$\frac{f_0^{n+1} - f_0^n}{\Delta t} = \partial_x(-b_0 f_0^n - a_0 f_1^n), \tag{8.3.10}$$

$$\frac{f_1^{n+1} - f_1^n}{\Delta t} = -a_1 \partial_x f_2^n - b_1 \partial_x f_0^n + E^* d_{10} f_0^n, \tag{8.3.11}$$

$$\frac{f_2^{n+1} - f_2^n}{\Delta t} = -a_2 \partial_x f_3^n - b_2 \partial_x f_2^n - a_1 \partial_x f_1^n + E^*(d_{20} f_0^n + d_{21} f_1^n). \tag{8.3.12}$$

The third moment f_3^n can be computed from the low-rank factors via

$$f_3^n = \sum_{ij} X_i^n S_{ij}^n \left\langle p_3(v) V_j^n \right\rangle_v. \tag{8.3.13}$$

Define $\mathcal{N}^n(x, v)$ by

$$\mathcal{N}^n(x, v) = w(v)[p_0(v) f_0^n(x) + p_1(v) f_1^n(x) + p_2(v) f_2^n(x)].$$

3. **K step:** Calculate $K_j^n(x) = \sum_i X_i^n(x) S_{ij}^n$, and perform a Forward Euler step of (8.3.4) to obtain

$K_j^{n+1}(x)$:

$$\frac{K_j^{n+1} - K_j^n}{\Delta t} = \left\langle V_j^n, P_\Phi^\perp D\left[E^*, \mathcal{N}^n, \sum_l K_l^n V_l^n\right]\right\rangle_{w^{-1}(v)}. \tag{8.3.14}$$

Perform a QR decomposition of $K_j^{n+1}(x)$ to obtain $X_i^{n+1}(x)$ and S'_{ij}.

4. **S step:** Perform a Forward Euler step of (8.3.5) to obtain S''_{ij}:

$$\frac{S''_{ij} - S'_{ij}}{\Delta t} = -\left\langle X_i^{n+1}V_j^n, P_\Phi^\perp D\left[E^*, \mathcal{N}^n, \sum_{kl} X_k^{n+1} S'_{kl} V_l^n\right]\right\rangle_{x,w^{-1}(v)}. \tag{8.3.15}$$

5. **L step:** Calculate $L_i^n(v) = \sum_j S''_{ij} V_j^n(v)$, and perform a Forward Euler step of (8.3.6) to obtain $L_i^{n+1}(v)$:

$$\frac{L_i^{n+1} - L_i^n}{\Delta t} = \left\langle X_i^{n+1}, P_\Phi^\perp D\left[E^*, \mathcal{N}^n, \sum_k X_k^{n+1} L_k^n\right]\right\rangle_x. \tag{8.3.16}$$

At this point, we must perform a QR decomposition of $L_i^{n+1}(v)$ to obtain the new velocity basis V_j^{n+1} and singular value matrix $S_{ij}^{n+1}(v)$. However, we also require that the velocity basis so obtained is orthogonal to Φ.

We accomplish this by prepending the three functions $w(v)\mathbf{p}(v)$ to the vector of functions $L_i^{n+1}(v)$. Denoting vector concatenation by square brackets, we compute a QR decomposition

$$[w(v)\mathbf{p}(v)\ \ V_j^{n+1}(v)]R = [w(v)\mathbf{p}(v)\ \ L_i^{n+1}(v)], \tag{8.3.17}$$

with respect to the function inner product $\langle\cdot\rangle_{w(v)^{-1}}$. The QR decomposition leaves the first three functions unchanged since they are already orthogonal, and guarantees that the resulting basis $V_j^{n+1}(v)$ is orthogonal to Φ, which is spanned by $w(v)\mathbf{p}(v)$.

The updated matrix of singular values is then given by the trailing $r \times r$ minor of R:

$$S_{ij}^{n+1} = R_{i+3,j+3}.$$

Proof of conservation

We now prove that our scheme satisfies the local conservation laws stated in Equations (8.1.6)-(8.1.8). A similar local conservation statement is satisfied by other conservative dynamical low-rank integrators such as [38, 35], and is important for ensuring that the solution has the properties and behavior of a local hyperbolic conservation law, such as finite wavespeeds. We formulate the discrete version of local conservation in the following theorem:

Theorem 8.3.1. *Define the conserved quantities of charge, momentum, and kinetic energy density at time level t^n as follows:*

$$\rho^n = \langle f^n, 1\rangle_v , \tag{8.3.18}$$

$$(\rho u)^n = \langle f^n, v\rangle_v , \tag{8.3.19}$$

$$\kappa^n = \left\langle f^n, \frac{1}{2}v^2\right\rangle_v . \tag{8.3.20}$$

The first-order integration algorithm of Section 8.3.2 satisfies the three local source-balance laws

$$\frac{\rho^{n+1}-\rho^n}{\Delta t} + \partial_x(\rho u)^n = 0, \tag{8.3.21}$$

$$\frac{(\rho u)^{n+1}-(\rho u)^n}{\Delta t} + 2\partial_x\kappa^n = \rho^n E^*, \tag{8.3.22}$$

$$\frac{\kappa^{n+1}-\kappa^n}{\Delta t} + \partial_x\left\langle f^n, \frac{v^3}{2}\right\rangle_v = (\rho u)^n E^*. \tag{8.3.23}$$

Proof. By (8.3.17), $P_\Phi V_j^n = 0$, so $\langle g^n, \boldsymbol{\phi}(v)\rangle_v = 0$ for all times n. Therefore the only contribution to the conserved quantities comes from f_0, f_1, and f_2.

To show mass conservation, we rewrite (8.3.10) using the recurrence and completeness properties of the polynomials $p_i(v)$, (8.2.6) and (8.2.8), to obtain

$$\begin{aligned}
\frac{\rho^{n+1}-\rho^n}{\Delta t} = q_{00}\frac{f_0^{n+1}-f_0^n}{\Delta t} &= q_{00}\nabla_{\boldsymbol{x}}(-b_0 f_0^n - a_0 f_1^n)\\
&= -q_{00}\nabla_{\boldsymbol{x}}\langle f^n, v p_0(v)\rangle_v\\
&= -\nabla_{\boldsymbol{x}}\langle f^n, v\rangle_v\\
&= -\nabla_{\boldsymbol{x}}(\rho u)^n.
\end{aligned}$$

This proves (8.3.21).

To show momentum conservation, we rewrite (8.3.11) in the same way:

$$\begin{aligned}
\frac{(\rho u)^{n+1}-(\rho u)^n}{\Delta t} &= q_{11}\frac{f_1^{n+1}-f_1^n}{\Delta t} + q_{10}\frac{f_0^{n+1}-f_0^n}{\Delta t}\\
&= q_{11}(-a_1\partial_x f_2^n - b_1\partial_x f_1^n - a_0\partial_x f_0^n + E^* d_{10} f_0^n) + q_{10}(-b_0\partial_x f_0^n - a_0\partial_x f_1^n)\\
&= -q_{11}\partial_x\langle f^n, v p_1(v)\rangle_v - q_{10}\partial_x\langle f^n, v p_0(v)\rangle_v + q_{11}E^* d_{10}\langle f^n, p_0(v)\rangle_v\\
&= -\partial_x\left\langle f^n, v^2\right\rangle_v + q_{11}E^*\langle f^n, p_1'(v)\rangle_v\\
&= -\partial_x\left\langle f^n, v^2\right\rangle_v - E^*\langle \partial_v f^n, v\rangle_v\\
&= -2\partial_x\kappa^n + E^*\rho^n.
\end{aligned}$$

For kinetic energy conservation, recall that $\kappa = \frac{1}{2}\left\langle |v|^2/2f\right\rangle_v$. Expanding (8.3.12) we obtain

$$\begin{aligned}
2\frac{\kappa^{n+1}-\kappa^n}{\Delta t} &= q_{22}\frac{f_2^{n+1}-f_2^n}{\Delta t} + q_{21}\frac{f_1^{n+1}-f_1^n}{\Delta t} + q_{20}\frac{f_0^{n+1}-f_0^n}{\Delta t} \\
&= -\partial_x \left\langle f^n, v[q_{22}p_2(v) + q_{21}p_1(v) + q_{20}p_0(v)]\right\rangle_v \\
&\quad + E^* \left\langle f^n, [q_{22}p_2'(v) + q_{21}p_1'(v)]\right\rangle_v \\
&= -\partial_x \left\langle f^n, v^3\right\rangle_v - E^* \left\langle \partial_v f^n, v^2\right\rangle_v .
\end{aligned}$$

Integrating by parts,

$$\frac{\kappa^{n+1}-\kappa^n}{\Delta t} + \partial_x \left\langle f^n, \frac{v^3}{2}\right\rangle_v = (\rho u)^n E^*.$$

This completes the proof. □

Theorem 8.3.1 shows that the microscopic current and kinetic energy satisfy the same source-balance laws as the full kinetic equation. The source terms in (8.3.22) and (8.3.23) reflect the exchange of momentum and energy between the particles and the electric field. For a physically accurate simulation, our discretization must not violate energy conservation during this exchange. that the total energy including electric field energy satisfies a local conservation law:

Corollary 8.3.1. *Define the total energy at time level* t^n *as*

$$e^n = \kappa^n + \frac{|E^n|^2}{2}. \tag{8.3.24}$$

Then the first-order integration algorithm of Section 8.3.2 with the choice of Ampère solve and $E^* = \frac{E^{n+1}+E^n}{2}$ *satisfies a local conservation law,*

$$\frac{e^{n+1}-e^n}{\Delta t} + \partial_x \left\langle f^n, \frac{v^3}{2}\right\rangle_v = 0. \tag{8.3.25}$$

Proof. The proof is a simple application of (8.3.23) and (8.3.9):

$$\begin{aligned}
\frac{e^{n+1}-e^n}{\Delta t} &= \frac{\kappa^{n+1}-\kappa^n}{\Delta t} + \frac{|E^{n+1}|^2 - |E^n|^2}{2\Delta t} \\
&= -\partial_x \left\langle f^n, \frac{v^3}{2}\right\rangle_v + \frac{E^{n+1}+E^n}{2}(\rho u)^n + \frac{E^{n+1}+E^n}{2}\frac{E^{n+1}-E^n}{\Delta t} \\
&= -\partial_x \left\langle f^n, \frac{v^3}{2}\right\rangle_v .
\end{aligned}$$

□

To summarize, the Ampère's law variation of the first-order integrator satisfies exact local conservation of charge and total energy, per (8.3.21) and (8.3.25) respectively. Momentum is not locally conserved since the mobile particles "push against" the background charge density ρ_0, whose equal and opposite momentum is neglected. However, a global conservation statement for momentum does hold in the special case of periodic boundary conditions for the Gauss's law variation of the first-order integrator, as shown in the following corollary.

Corollary 8.3.2. *If Gauss's law and the uncentered electric field* $E^* = E^n$ *are chosen in the first-order integrator, then*

$$\int_{\Omega_x} \frac{(\rho u)^{n+1} - (\rho u)^n}{\Delta t} \, \mathrm{d}x = 0.$$

Proof. Integrate (8.3.22) over the periodic domain Ω_x to obtain

$$\int_{\Omega_x} \frac{(\rho u)^{n+1} - (\rho u)^n}{\Delta t} \, \mathrm{d}x = \int_{\Omega_x} \rho^n E^* \, \mathrm{d}x = -\int_{\Omega_x} \rho^n \partial_x \varphi^n \, \mathrm{d}x = \int_{\Omega_x} (\partial_x^2 \varphi^n - \rho_0) \partial_x \varphi^n \, \mathrm{d}x = 0.$$

We have used (8.3.7) and (8.3.8), and eliminated all integrals of total derivatives in x. □

Fully discrete conservation

The conservation statements encapsulated by Theorem 8.3.1 and Corollaries 8.3.1 and 8.3.2 are semi-discrete. Up to this point we have discretized the time and velocity variables but have left the physical space variable x continuous. It is not a challenge to translate the semi-discrete conservation properties we have shown so far into fully discrete conservation. This can be done with any standard conservative discretization of the ∂_x operator appearing in (8.3.10), (8.3.11), and (8.3.12). Examples include any discretization that can be cast in a flux-differencing form [78].

Furthermore, a fully discrete scheme using the Gauss's law variation of the first-order integrator can be designed to exactly conserve momentum on a periodic domain. This requires that the discrete ∂_x operators of (8.3.7) and (8.3.8) satisfy a summation-by-parts identity. One such discretization is the ubiquitous second-order centered finite difference scheme with stencil $[1, -2, 1]$.

8.3.3 Second-order time integrator

The projector-splitting framework may be formally extended to second-order accuracy by using a Strang splitting of (8.3.3). A Strang splitting of the dynamical low-rank projection has been used in [36] and [32] to obtain second-order accurate solutions to certain problems. We stress that it has not been proven that the Strang splitting of the DLR projection is robust to vanishing singular values, as it has for the first-order splitting [70]. That is, no proof of robust second-order accuracy exists. Nevertheless, the second-order Strang splitting scheme shows significant practical benefits on plasma problems, as demonstrated for example in [32]. To achieve overall second-order accuracy, some care is required when coupling the DLR scheme with the time splitting for $\mathcal{N}$ and E. Here we present one such scheme and prove that it exactly conserves energy.

1. **Half step of conserved quantities**: $(f_0^n, f_1^n, f_2^n) \mapsto (f_0^{n+1/2}, f_1^{n+1/2}, f_2^{n+1/2})$ using g^n, E^n:

$$\frac{f_0^{n+1/2} - f_0^n}{\Delta t/2} = \partial_x(-b_0 f_0^n - a_0 f_1^n), \tag{8.3.26}$$

$$\frac{f_1^{n+1/2} - f_1^n}{\Delta t/2} = \partial_x(-a_1 f_2^n - b_1 f_1^n - a_0 f_0^n) + E^n d_{10} f_0^n, \tag{8.3.27}$$

$$\frac{f_2^{n+1/2} - f_2^n}{\Delta t/2} = \partial_x(-a_2 f_3^n - b_2 f_2^n - a_1 f_1^n) + E^n (d_{20} f_0^n + d_{21} f_1^n), \tag{8.3.28}$$

where f_3^n is defined as in (8.3.13). Define

$$\mathcal{N}^{n+1/2} = w(v)[p_0(v) f_0^{n+1/2} + p_1(v) f_1^{n+1/2} + p_2(v) f_2^{n+1/2}].$$

2. **Ampère solve**: $E_n \mapsto E^{n+1}$ using $f_0^{n+1/2}, f_1^{n+1/2}$:

$$\frac{E^{n+1} - E^n}{\Delta t} = -J^{n+1/2}, \quad J^{n+1/2} = q_{10} f_0^{n+1/2} + q_{11} f_1^{n+1/2}. \tag{8.3.29}$$

Define $E^{n+1/2} = \frac{E^{n+1}+E^n}{2}$.

3. **K step**: $(X^n, S^n, V^n) \mapsto (X^{n+1/2}, S^1, V^n)$ using $\mathcal{N}^{n+1/2}, E^{n+1/2}$.
4. **S step**: $(X^{n+1/2}, S^1, V^n) \mapsto (X^{n+1/2}, S^2, V^n)$ using $\mathcal{N}^{n+1/2}, E^{n+1/2}$.
5. **L step**: $(X^{n+1/2}, S^2, V^n) \mapsto (X^{n+1/2}, S^{n+1/2}, V^{n+1/2})$ using $\mathcal{N}^{n+1/2}, E^{n+1/2}$.
6. **L step**: $(X^{n+1/2}, S^{n+1/2}, V^{n+1/2}) \mapsto (X^{n+1/2}, S^3, V^{n+1})$ using $\mathcal{N}^{n+1/2}, E^{n+1/2}$.
7. **S step**: $(X^{n+1/2}, S^3, V^{n+1}) \mapsto (X^{n+1/2} S^4, V^{n+1})$ using $\mathcal{N}^{n+1/2}, E^{n+1/2}$.
8. **K step**: $(X^{n+1/2}, S^4, V^{n+1}) \mapsto (X^{n+1}, S^{n+1}, V^{n+1})$ using $\mathcal{N}^{n+1/2}, E^{n+1/2}$.
9. $(f_0^n, f_1^n, f_2^n) \mapsto (f_0^{n+1}, f_1^{n+1}, f_2^{n+1})$ using $g^{n+1/2}, E^{n+1/2}$, and $(f_0, f_1, f_2)^{n+1/2}$:

$$\frac{f_0^{n+1} - f_0^n}{\Delta t} = \partial_x(-b_0 f_0^{n+1/2} - a_0 f_1^{n+1/2}), \tag{8.3.30}$$

$$\frac{f_1^{n+1} - f_1^n}{\Delta t} = \partial_x(-a_1 f_2^{n+1/2} - b_1 f_1^{n+1/2} - a_0 f_0^{n+1/2}) + E^{n+1/2} d_{10} f_0^{n+1/2}, \tag{8.3.31}$$

$$\frac{f_2^{n+1} - f_2^n}{\Delta t} = \partial_x(-a_2 f_3^{n+1/2} - b_2 f_2^{n+1/2} - a_1 f_1^{n+1/2}) + E^{n+1/2}(d_{20} f_0^{n+1/2} + d_{21} f_1^{n+1/2}), \tag{8.3.32}$$

where

$$f_3^{n+1/2} = \sum_{ij} X_i^{n+1/2} S_{ij}^{n+1/2} \left\langle p_3(v) V_j^{n+1/2} \right\rangle_v .$$

Note that in step 9 we make use of $g^{n+1/2}$, which is defined in terms of $(X^{n+1/2}, S^{n+1/2}, V^{n+1/2})$. This means that we cannot easily combine steps 5 and 6 into a single substep of size Δt as is common in Strang splitting schemes. To achieve overall second-order accuracy in time, each of the low-rank factor steps 3-8 must be accomplished using a standard time integration scheme that is at least second-order accurate. We choose the SSPRK2 scheme [45], which for an autonomous ordinary differential equation $q'(t) = F(q)$ is

$$\begin{aligned} q^* &= q^n + \Delta t F(q^n) \\ q^{n+1} &= \frac{q^n}{2} + \frac{1}{2}\left(q^* + \Delta t F(q^*)\right). \end{aligned} \tag{8.3.33}$$

Proof of charge and energy conservation

Theorem 8.3.2. *The second-order integrator satisfies local conservation of charge and total energy:*

$$\frac{\rho^{n+1} - \rho^n}{\Delta t} + \partial_x J^{n+1/2} = 0, \tag{8.3.34}$$

$$\frac{e^{n+1} - e^n}{\Delta t} + \partial_x \left\langle \frac{v^3}{2}, f^{n+1/2} \right\rangle_v = 0, \tag{8.3.35}$$

where ρ^n *and* e^n *are defined as in* (8.3.18) *and* (8.3.24) *respectively, and* $J^{n+1/2}$ *is defined as in* (8.3.19) *but at* $t^{n+1/2}$.

Proof. By (8.3.17), $P_\Phi V_j^n = \mathbf{0}$, so $\langle g^n, \boldsymbol{\phi}(v) \rangle_v = 0$ for all times n. Therefore, the only contributions to the charge and total energy come from f_0, f_1, f_2 and E. From equations (8.3.30) and (8.3.32), by following the proof of Theorem 8.3.1 with fluxes and source terms evaluated at $t^{n+1/2}$, we obtain

$$\frac{\rho^{n+1} - \rho^n}{\Delta t} + \partial_x J^{n+1/2} = 0, \tag{8.3.36}$$

$$\frac{\kappa^{n+1} - \kappa^n}{\Delta t} + \partial_x \left\langle \frac{v^3}{2}, f^{n+1/2} \right\rangle_v = J^{n+1/2} E^{n+1/2}. \tag{8.3.37}$$

Combining (8.3.37) with (8.3.29) and the definition of $E^{n+1/2}$, we derive the stated total energy conservation law (8.3.35). □

Corollary 8.3.3. *A fully discrete scheme for the second-order integrator of Section 8.3.3, which uses a conservative spatial discretization for the* ∂_x *operators appearing in equations* (8.3.26)-(8.3.28) *and* (8.3.30)-(8.3.32) *will satisfy exact discrete charge and energy conservation.*

8.3.4 Substeps for low-rank factors

In this section we expand each of the low-rank factors' equation of motion, (8.3.4), (8.3.5), (8.3.6). Our purpose is to demonstrate that the proposed algorithm is efficient in the sense of not requiring operations that have a computational cost of $\mathcal{O}(N_x N_v)$, where N_x, N_v represent the degrees of freedom used to discretize x, v respectively. To avoid proliferation of indices, in this section we consider only the case of simple Forward Euler steps used in the first-order integrator. The computational cost of the second-order integrator, whose substeps are themselves SSPRK2 steps, differs by only a constant factor.

The terms on the right-hand side of the K and S steps can be simplified by the following observation.

Proposition 8.3.1. *Let the basis functions V_j^n be evolved according to the first-order time integrator of Section 8.3.2. Then, for all $h(v) \in L^2(\Omega_v, w^{-1}(v))$, the following identity holds for all $V_j^n(v)$:*

$$\left\langle V_j^n(v), P_\Phi^\perp h(v) \right\rangle_{w^{-1}(v)} = \left\langle V_j^n(v), h(v) \right\rangle_{w^{-1}(v)} . \tag{8.3.38}$$

Proof. Decompose $h(v)$ as $h(v) = P_\Phi h(v) + P_\Phi^\perp h(v)$, and use $\left\langle V_j^n(v), P_\Phi h(v) \right\rangle_{w^{-1}(v)} = 0$, since V_j^n is orthogonal to the range of P_Φ. ☐

Applying boundary conditions in v

For a bounded velocity domain such as the finite difference discretization we will discuss below, we must apply a boundary condition in v. This is clear in the case of the L step, where the boundary condition is applied to a hyperbolic term. It is also true for the K and S steps, where the boundary condition is required for the evaluation of a derivative under an integral. In [62] the authors show how to treat boundary conditions in x in the dynamical low-rank framework; we use the same approach. The idea is that the Dirichlet boundary condition $g(x, v_b, t) = -\mathcal{N}(x, v_b, t)$ is not imposed on the basis functions V_j directly. Rather, by projecting onto X_i, we can see that the boundary condition should be applied to the weighted function basis L_i:

$$\langle X_i, g(x, v_b, t) \rangle_x = L_i(v_b, t) = \langle X_i, -\mathcal{N}(x, v_b, t) \rangle_x .$$

Then, when expanding the low-rank equations of motion, we interpret terms involving $\partial_v g$ as $\sum_k X_k \partial_v L_k$, rather than the usual $\sum_{kl} X_k S_{kl} \partial_v V_l$, with the plan of applying the boundary condition on L_k in the course of evaluating the derivative.

Our assumption of periodic boundary conditions in x saves us the trouble of performing the same transformation to terms involving $\partial_x g$; however this poses no essential difficulty, and nontrivial nonperiodic boundary conditions in x are a straightforward extension of this scheme.

K step

We may expand (8.3.14) by substituting (8.3.1) into (8.3.4). Since the collisional moments u and T are held constant during the K step, at a time level t^n, we may write

$$C(\mathcal{N}^n + g) = C^n(\mathcal{N}^n + g) = C^n(\mathcal{N}^n) + C^n(g),$$

where

$$C^n(f) = \nu\partial_v(T^n\partial_v f + (v - u^n)f).$$

After using Proposition 8.3.1 to eliminate appearances of $P_\Phi^\perp$, this gives

$$\frac{K_j^{n+1} - K_j^n}{\Delta t} = \left\langle V_j^n, D\left[E^*, \mathcal{N}^n, \sum_l K_l^n V_l^n\right]\right\rangle_{w^{-1}(v)} \tag{8.3.39}$$

$$\begin{aligned} &= -\left\langle V_j^n, v\nabla_{\boldsymbol{x}}\mathcal{N}^n + E^*\nabla_{\boldsymbol{v}}\mathcal{N}^n\right\rangle_{w^{-1}(v)} + \left\langle V_j^n, C^n(\mathcal{N}^n)\right\rangle_{w^{-1}(v)} \\ &\quad - \sum_l \left\langle V_j^n, vV_l^n\right\rangle_{w^{-1}(v)} \nabla_{\boldsymbol{x}} K_l^n - \sum_k X_k^n \left\langle V_j^n, \nabla_{\boldsymbol{v}} L_k^n\right\rangle_{w^{-1}(v)} E^* \\ &\quad + \nu\sum_k X_k^n \left[T^n \left\langle V_j^n, \nabla_{\boldsymbol{v}}^2 L_k^n\right\rangle_{w^{-1}(v)} + \left\langle V_j^n, \nabla_{\boldsymbol{v}}(vL_k^n)\right\rangle_{w^{-1}(v)} - u^n \left\langle V_j^n, \nabla_{\boldsymbol{v}} L_k^n\right\rangle_{w^{-1}(v)}\right], \end{aligned} \tag{8.3.40}$$

where $L_i^n = \sum_j S_{ij}^n V_j^n$. Derivative terms involving L_i^n are to be calculated using the projected boundary conditions on L:

$$L_i^n(v_b) = \left\langle X_i^n, -\mathcal{N}^n(x, v_b)\right\rangle_x .$$

Because some of the inner products appearing in (8.3.39) will also appear in the S step, we can save some computational effort by computing and saving the vectors and matrices in (8.3.39):

$$\boldsymbol{y}_j^1 = \left\langle V_j^n, v\nabla_{\boldsymbol{x}}\mathcal{N}^n\right\rangle_{w^{-1}(v)}, \quad \boldsymbol{y}_j^2 = \left\langle V_j^n, E^*\nabla_{\boldsymbol{v}}\mathcal{N}^n\right\rangle_{w^{-1}(v)}, \quad \boldsymbol{y}_j^3(x) = \left\langle V_j^n, C^n(\mathcal{N}^n)\right\rangle_{w^{-1}(v)} \tag{8.3.41}$$

$$\boldsymbol{A}_{jl} = \left\langle V_j^n, vV_l^n\right\rangle_{w^{-1}(v)}, \quad \boldsymbol{D}_{jk}^1 = \left\langle V_j^n, \partial_v L_k^n\right\rangle_{w^{-1}(v)} \tag{8.3.42}$$

$$\boldsymbol{D}_{jk}^2 = \left\langle V_j^n, \partial_v^2 L_k^n\right\rangle_{w^{-1}(v)}, \quad \boldsymbol{G}_{jk} = \left\langle V_j^n, \partial_v(vL_k^n)\right\rangle_{w^{-1}(v)} \tag{8.3.43}$$

Cost: $\mathcal{O}(3r(N_v + N_x) + r^2 N_v)$, by taking advantage of the rank-3 structure of $\mathcal{N}$.

Then the Forward Euler step for K can be written concisely as

$$\frac{K_j^{n+1} - K_j^n}{\Delta t} = -\boldsymbol{y}_j^1 - \boldsymbol{y}_j^2 + \boldsymbol{y}_j^3 - \sum_l \boldsymbol{A}_{jl} \partial_x K_l^n + \sum_k X_k^n \left[-E^* \boldsymbol{D}_{jk}^1 + \nu (T^n \boldsymbol{D}_{jk}^2 + \boldsymbol{G}_{jk} - u^n \boldsymbol{D}_{jk}^1) \right].$$

Cost: $\mathcal{O}(r^2 N_x)$.

Remark 4. Additional care must be taken in the implementation of the second-order integrator to ensure each of the intermediate vectors and matrices in (8.3.41) is defined and computed in terms of quantities with the correct time levels on the right-hand side.

S step

Because the X_i basis has been updated in the K step, we must re-project the boundary conditions in v onto X_i^{n+1}. Define $\tilde{L}_i = \sum_j S'_{ij} V_j^n$, then the projected boundary conditions are

$$\tilde{L}_i(v_b) = \left\langle X_i^{n+1}, -\mathcal{N}^n(x, v_b) \right\rangle_x .$$

To expand (8.3.15), substitute (8.3.1) into (8.3.15) to obtain

$$\begin{aligned}
\frac{S''_{ij} - S'_{ij}}{\Delta t} &= -\left\langle X_i^{n+1} V_j^n, D\left[E^*, \mathcal{N}^n, \sum_{kl} X_k^{n+1} S'_{kl} V_l^n \right] \right\rangle_{x, w^{-1}(v)} \qquad (8.3.44) \\
&= \left\langle X_i^{n+1} V_j^n, v \nabla_{\boldsymbol{x}} \mathcal{N}^n + E^* \nabla_{\boldsymbol{v}} \mathcal{N}^n \right\rangle_{x, w^{-1}(v)} - \left\langle X_i^{n+1} V_j^n, C^n(\mathcal{N}^n) \right\rangle_{x, w^{-1}(v)} \\
&\quad + \sum_{kl} S'_{kl} \left\langle X_i^{n+1}, \nabla_{\boldsymbol{x}} X_k^{n+1} \right\rangle_x \left\langle V_j^n, v V_l^n \right\rangle_{w^{-1}(v)} + \sum_k \left\langle X_i^{n+1}, E^* X_k^{n+1} \right\rangle_x \left\langle V_j^n, \nabla_{\boldsymbol{v}} \tilde{L}_k \right\rangle_{w^{-1}(v)} \\
&\quad - \nu \sum_k \left[\left\langle X_i^{n+1}, T^n X_k^{n+1} \right\rangle_x \left\langle V_j^n, \nabla_{\boldsymbol{v}}^2 \tilde{L}_k \right\rangle_{w^{-1}(v)} \right. \qquad (8.3.45) \\
&\quad + \left\langle X_i^{n+1}, X_k^{n+1} \right\rangle_x \left\langle V_j^n, \nabla_{\boldsymbol{v}} (v \tilde{L}_k) \right\rangle_{w^{-1}(v)} \\
&\quad \left. - \left\langle X_i^{n+1}, u^n X_k^{n+1} \right\rangle_x \left\langle V_j^n, \nabla_{\boldsymbol{v}} \tilde{L}_k \right\rangle_{w^{-1}(v)} \right]
\end{aligned}$$

We can save computational effort again by precomputing certain matrices which will appear in the L step. First, compute

$$\tilde{\boldsymbol{D}}_{jk}^1 = \left\langle V_j^n, \partial_v \tilde{L}_k \right\rangle_{w^{-1}(v)} \qquad \tilde{\boldsymbol{D}}_{jk}^2 = \left\langle V_j^n, \partial_v^2 \tilde{L}_k \right\rangle_{w^{-1}(v)}, \qquad \tilde{\boldsymbol{G}}_{jk} = \left\langle V_j^n, \partial_v (v \tilde{L}_k) \right\rangle_{w^{-1}(v)} .$$

Cost: $\mathcal{O}(r^2 N_v)$.

Then take advantage of the rank-3 structure of $\mathcal{N}$ to compute the inner products in both x and v:

$$\boldsymbol{Z}_{ij}^1 = \left\langle X_i^{n+1} V_j^n, v \nabla_{\boldsymbol{x}} \mathcal{N}^n \right\rangle_{x, w^{-1}(v)}, \quad \boldsymbol{Z}_{ij}^2 = \left\langle X_i^{n+1} V_j^n, E^* \nabla_{\boldsymbol{v}} \mathcal{N}^n \right\rangle_{x, w^{-1}(v)},$$

$$\boldsymbol{Z}^3_{ij} = \left\langle X_i^{n+1} V_j^n, C^n(\mathcal{N}^n) \right\rangle_{x,w^{-1}(v)} .$$

Cost: $\mathcal{O}(3r^2(N_v + N_x))$.

Finally, compute the matrices from inner products in x:

$$\boldsymbol{B}_{ik} = \left\langle X_i^{n+1}, \partial_x X_k^{n+1} \right\rangle_x , \quad \boldsymbol{F}_{ik} = \left\langle X^{n+1}, E^* X_k^{n+1} \right\rangle_x$$

$$\boldsymbol{N}_{ik} = \left\langle X_i^{n+1}, T^n X_k^{n+1} \right\rangle_x , \quad \delta_{ik} = \left\langle X^{n+1}, X_k^{n+1} \right\rangle_x , \quad \boldsymbol{Q}_{ik} = \left\langle X^{n+1}, u^n X_k^{n+1} \right\rangle_x .$$

Cost: $\mathcal{O}(r^2 N_x)$.

Now the Forward Euler step for S can be written concisely as

$$\frac{S''_{ij} - S'_{ij}}{\Delta t} = \boldsymbol{Z}^1_{ij} + \boldsymbol{Z}^2_{ij} - \boldsymbol{Z}^3_{ij} + \sum_{kl} S'_{kl} \boldsymbol{B}_{ik} \boldsymbol{A}_{jl} - \sum_k \left[-\boldsymbol{F}_{ik} \tilde{\boldsymbol{D}}^1_{jk} + \nu(\boldsymbol{N}_{ik} \tilde{\boldsymbol{D}}^2_{jk} + \delta_{ik} \tilde{\boldsymbol{G}}_{jk} - \boldsymbol{Q}_{ik} \tilde{\boldsymbol{D}}^1_{jk}) \right]$$

Cost: $\mathcal{O}(r^4)$.

L step

The L step is calculated similarly to the K and S steps. Because $P_\Phi^\perp$ involves projecting only in v, it commutes with the x inner product $\langle \cdot \rangle_x$, which lets us pull the projection $P_\Phi^\perp$ out of each low-rank projection. Recall that $P_\Phi^\perp = I - P_\Phi$ where P_Φ is defined in (8.2.10) Plug (8.3.1) into (8.3.6) to obtain

$$\frac{L_i^{n+1} - L_i^n}{\Delta t} = \left\langle X_i^{n+1}, P_\Phi^\perp D \left[E^*, \mathcal{N}^n, \sum_k X_k^{n+1} L_k^n \right] \right\rangle_x \tag{8.3.46}$$

$$\begin{aligned} &= -P_\Phi^\perp \left\langle X_i^{n+1}, (v \nabla_{\boldsymbol{x}} \mathcal{N}^n + E^* \nabla_{\boldsymbol{v}} \mathcal{N}^n) \right\rangle_x + P_\Phi^\perp \left\langle X_i^{n+1}, C^n(\mathcal{N}^n) \right\rangle_x \\ &\quad - P_\Phi^\perp \left(\sum_k v \left\langle X_i^{n+1}, \nabla_{\boldsymbol{x}} X_k^{n+1} \right\rangle_x L_k^n + \sum_k \left\langle X_i^{n+1}, E^* X_k^{n+1} \right\rangle_x \nabla_{\boldsymbol{v}} L_k^n \right) \\ &\quad + P_\Phi^\perp \nu \sum_k \left[\left\langle X_i^{n+1}, T^n X_k^{n+1} \right\rangle_x \nabla_{\boldsymbol{v}}^2 L_k^n + \left\langle X_i^{n+1}, X_k^{n+1} \right\rangle_x \nabla_{\boldsymbol{v}} (v L_k^n) - \left\langle X_i^{n+1}, u^n X_k^{n+1} \right\rangle_x \nabla_{\boldsymbol{v}} L_k^n \right] . \end{aligned} \tag{8.3.47}$$

Compute vectors appearing in (8.3.46):

$$\boldsymbol{z}_i^1(v) = \left\langle X_i^{n+1}, v \nabla_{\boldsymbol{x}} \mathcal{N}^n \right\rangle_x , \quad \boldsymbol{z}_i^2(v) = \left\langle X_i^{n+1}, E^* \nabla_{\boldsymbol{v}} \mathcal{N}^n \right\rangle_x , \quad \boldsymbol{z}_i^3(v) = \left\langle X_i, C(\mathcal{N}^n; U) \right\rangle_x .$$

Cost: $\mathcal{O}(3r(N_v + N_x)$.

The Forward Euler step for L can be written concisely as

$$\frac{L_i^{n+1} - L_i^n}{\Delta t} = P_\Phi^\perp(-\boldsymbol{z}_i^1(v) - \boldsymbol{z}_i^2(v) + \boldsymbol{z}_i^3(v)) \tag{8.3.48}$$

$$+ P_\Phi^\perp\left[\sum_k (-vL_k^n \boldsymbol{B}_{ik} - \boldsymbol{F}_{ik}\partial_v L_k^n + \nu(\boldsymbol{N}_{ik}\partial_v^2 L_k^n + \delta_{ik}\partial_v(vL_k^n) - \boldsymbol{Q}_{ik}\partial_v L_k^n))\right]. \tag{8.3.49}$$

Cost: $\mathcal{O}(r^2 N_v)$.

The total computational cost of all three substeps scales as $\mathcal{O}(r^2(N_x + N_v) + r^4)$. The asymptotic advantages do not begin to tell in one dimension, but in 2D2V and 3D3V kinetic applications, we can easily have $N_x, N_v \gg r^2$, in which case the efficiency gains from dynamical low-rank approximation are quite substantial.

8.4 Spatial discretization

In this section we discuss the discretization of the equations of Section 8.3 in x. We will continue to leave the v-discretization unspecified for now, continuing our formulation in terms of continuous inner products in v. In x, we use a simple finite difference scheme on an equispaced grid, with grid separation denoted Δx, and points x_i. Functions are approximated directly on the grid, so that $u(x_i) \approx u_i$. For discretization of the inner product $\langle\cdot\rangle_x$ over physical space, we use the Trapezoidal rule. On periodic domains, the Trapezoidal rule has spectral convergence, while for nonperiodic domains, its second-order convergence is sufficient and matches the order of accuracy of our finite difference discretization of the hyperbolic terms, described below. The Poisson equation is solved using a standard centered finite difference stencil, also of second order accuracy. Finally, for non-hyperbolic first-order derivatives appearing inside of an inner product in x, we use a second-order centered finite difference approximation to ∂_x. Discussion of our hyperbolic finite difference discretization is below.

8.4.1 Finite difference discretization of hyperbolic terms

Because of the complexity of our time discretization, additional care is required to ensure that the overall hyperbolic structure of the kinetic equation is preserved in our discretization of x and v. In particular, four of the equations described in Section 8.3 have an advective form, involving the x-derivative of one or more unknowns. Ignoring the other (source) terms in those equations, they are:

$$\partial_t \begin{pmatrix} f_0 \\ f_1 \\ f_2 \end{pmatrix} + \begin{pmatrix} b_0 & a_0 & 0 \\ a_0 & b_1 & a_1 \\ 0 & a_1 & b_2 \end{pmatrix} \partial_x \begin{pmatrix} f_0 \\ f_1 \\ f_2 \end{pmatrix} + \partial_x \begin{pmatrix} 0 \\ 0 \\ a_2 f_3 \end{pmatrix} = RHS, \tag{8.4.1}$$

and

$$\partial_t K_j + \langle V_j, v\partial_x \mathcal{N}\rangle_{w^{-1}(v)} + \sum_l \boldsymbol{A}_{jl}\partial_x K_l = RHS, \tag{8.4.2}$$

where $\boldsymbol{A}_{jl}$ is defined according to (8.3.42).

Proposition 8.4.1. *Equations* (8.4.1) *and* (8.4.2) *together form a globally hyperbolic system of partial differential equations.*

Proof. Note that, by (8.2.6),

$$\langle V_j, v\partial_x \mathcal{N}\rangle_{w^{-1}(v)} = \langle V_j, a_2\partial_x f_2(x,t)w(v)p_3(v)\rangle_{w^{-1}(v)} = \langle V_j, w(v)p_3(v)\rangle_{w^{-1}(v)} a_2\partial_x f_2(x,t),$$

where we have used Proposition 8.3.1 to eliminate terms in the span of $w(v)\mathbf{p}(v)$ inside the inner product. If we abbreviate the inner product by $\langle V_j, w(v)p_3(v)\rangle_{w^{-1}(v)} = \hat{\boldsymbol{y}}_j$, we can write the combined system in quasilinear form like so:

$$\partial_t \begin{pmatrix} f_0 \\ f_1 \\ f_2 \\ K_1 \\ K_2 \\ \vdots \\ K_r \end{pmatrix} + \underbrace{\begin{pmatrix} b_0 & a_0 & & & & & \\ a_0 & b_1 & a_1 & & & & \\ & a_1 & b_2 & a_2q_1 & a_2q_2 & \dots & a_2q_r \\ & & a_2\hat{\boldsymbol{y}}_1 & & & & \\ & & a_2\hat{\boldsymbol{y}}_2 & & \boldsymbol{A}_{jl} & & \\ & & \vdots & & & & \\ & & a_2\hat{\boldsymbol{y}}_r & & & & \end{pmatrix}}_{\boldsymbol{A}_0} \partial_x \begin{pmatrix} f_0 \\ f_1 \\ f_2 \\ K_1 \\ K_2 \\ \vdots \\ K_r \end{pmatrix} = RHS. \tag{8.4.3}$$

Here we have rewritten $f_3(x)$ as a linear combination of the $K_j(x)$, with coefficients q_j. Per (8.3.13), we have $q_j = \langle p_3(v), V_j\rangle_v$. But then

$$\hat{\boldsymbol{y}}_j = \langle V_j, w(v)p_3(v)\rangle_{w^{-1}(v)} = \langle V_j, p_3(v)\rangle_v = q_j.$$

Additionally, it is clear from its definition that $\boldsymbol{A}_{jl}$ is symmetric. Therefore, the matrix $\boldsymbol{A}_0$ is symmetric, and thus is strictly hyperbolic with real eigenvalues. □

The hyperbolicity of the overall system gives us confidence in applying standard discretization techniques for hyperbolic conservation laws. As a spatial discretization, we opt for a Shu-Osher conservative finite difference discretization based on upwind flux splitting. The hyperbolic system is always upwinded "all at once", despite the fact that we solve the first three and the trailing r equations separately.

To be precise, fix V_j and let $R\Lambda R^{-1} = \boldsymbol{A}_0$ be the eigendecomposition of $\boldsymbol{A}_0(V_j)$. Denote by $\boldsymbol{w}$ the length-$r+3$ vector $\boldsymbol{w} = [f_0, f_1, f_2, K_1, \ldots, K_r]^T$. Let P_1 and P_2 be the projection matrices consisting of the first three and trailing r rows of an $(r+3) \times (r+3)$ identity matrix; they select the rows of $\boldsymbol{w}$ corresponding to f_0, f_1, f_2 and K_j respectively:

$$P_1\boldsymbol{w} = [f_0, f_1, f_2]^T, \quad P_2\boldsymbol{w} = [K_1, K_2, \ldots, K_r]^T. \tag{8.4.4}$$

The flux terms appearing in (8.4.1) can be written as

$$\begin{pmatrix} b_0 & a_0 & 0 \\ a_0 & b_1 & a_1 \\ 0 & a_1 & b_2 \end{pmatrix} \partial_x \begin{pmatrix} f_0 \\ f_1 \\ f_2 \end{pmatrix} + \partial_x \begin{pmatrix} 0 \\ 0 \\ a_2 f_3 \end{pmatrix} = P_1 \partial_x \boldsymbol{A}_0 \boldsymbol{w}, \tag{8.4.5}$$

and the flux terms in (8.4.2) as

$$\langle V_j, v\partial_x \mathcal{N}\rangle_{w^{-1}(v)} + \sum_l \boldsymbol{A}_{jl} \partial_x K_l = P_2 \partial_x \boldsymbol{A}_0 \boldsymbol{w}. \tag{8.4.6}$$

The combined flux of f_0, f_1, f_2, and K_j is $F(\boldsymbol{w}) = \boldsymbol{A}_0\boldsymbol{w}$. Notionally, we approximate $\partial_x F(\boldsymbol{w})$ by a conservative flux difference:

$$\partial_x F(\boldsymbol{w}^n) = \frac{1}{\Delta x}(\hat{F}_{i+1/2} - \hat{F}_{i-1/2}),$$

where $\hat{F}_{i+1/2}$ is the numerical flux between cells i and $i+1$. The numerical flux is split and upwinded according to the eigendecomposition of $\boldsymbol{A}_0$. For example, a first-order upwind finite difference scheme would use

$$\hat{F}_{i+1/2} = \hat{F}^-_{i+1/2} + \hat{F}^+_{i+1/2} = R\Lambda^- R^{-1}\boldsymbol{w}_{i+1} + R\Lambda^+ R^{-1}\boldsymbol{w}_i. \tag{8.4.7}$$

In this work we use higher-order spatial reconstructions of the split flux, following the conservative finite difference framework of [91]. Except where noted, we use a MUSCL reconstruction [104] with the Monotonized-Central limiter [103] which is second-order in space. We then approximate the hyperbolic terms in (8.4.1) and (8.4.2) as

$$P_s \partial_x F(\boldsymbol{w}) = P_s \frac{1}{\Delta x}\left(\hat{F}_{i+1/2} - \hat{F}_{i-1/2}\right) \tag{8.4.8}$$

for $P_s \in \{P_1, P_2\}$. Importantly, we apply the projection P_s *after* the flux splitting and upwind reconstruction. We base the upwinding procedure on the whole size $r+3$ eigenvector decomposition of $\boldsymbol{w}$, rather than upwinding systems of size 3 and size r separately.

8.5 Two velocity space discretizations

Up to this point we have discretized in time and in physical space, but not yet in velocity space. The macro-micro decomposition described in Section 8.2 is made with reference to a particular set of orthogonal polynomials in velocity. The first three of these result in an unclosed system of three moment equations, (8.2.13). However, the remaining microscopic equations, which are projected into the dynamical low-rank equations of motion (8.3.39), (8.3.44), and (8.3.46), are still continuous in velocity and must be discretized. One of the key benefits of our method is that it is generic over the choice of velocity space discretization. To illustrate this, we present a pair of velocity space discretizations based on orthogonal polynomial expansions, for which the polynomials p_0, p_1, p_2 appear as the first three polynomials in one of the classical orthogonal polynomial families.

The first discretization we will discuss is a global Hermite spectral method, for which p_0, p_1, p_2 naturally enter as the first three Hermite polynomials. This method lets us discretize an unbounded velocity domain without arbitrary truncation. The second discretization uses a truncated velocity space, with p_0, p_1, p_2 chosen as the first three scaled Legendre polynomials. The resulting equations for the non-conserved part g are solved with a standard upwind finite difference method. Our aim with this pair of discretizations is to demonstrate the flexibility of the underlying macro-micro decomposition, which may be combined with whatever velocity space discretization is most convenient for the problem at hand.

8.5.1 Asymmetrically-weighted Hermite spectral method

The first velocity space discretization we use is the one-dimensional variant of the Hermite spectral discretization described in Section 4.2. To put this discretization in the notation of this chapter, the Hermite polynomials are the orthogonal polynomial family defined on $\Omega_v = \mathbb{R}$ with weight function

$$w(v) = \frac{1}{v_0\sqrt{2\pi}} e^{-\frac{v^2}{2v_0^2}},$$

where v_0 is a reference velocity which sets the width of the basis. The Hermite polynomials satisfy the following orthogonality relation:

$$\int_{\Omega_v} w(v) He_n\left(\frac{v}{v_0}\right) He_m\left(\frac{v}{v_0}\right) \mathrm{d}v = \delta_{nm}.$$

Their three-term recurrence relation, given in matrix form in (4.2.5), has coefficients $a_n = v_0\sqrt{n+1}$ and $b_n = 0$. The first several Hermite polynomials are

$$
\begin{aligned}
p_0(v) &= He_0\left(\frac{v}{v_0}\right) = 1 & \qquad p_2(v) &= He_2\left(\frac{v}{v_0}\right) = \frac{(v/v_0)^2 - 1}{\sqrt{2}} \\
p_1(v) &= He_1\left(\frac{v}{v_0}\right) = v/v_0 & \qquad p_3(v) &= He_3\left(\frac{v}{v_0}\right) = \frac{(v/v_0)^3 - v/v_0}{\sqrt{6}}.
\end{aligned}
$$

From these definitions, it is simple to verify the identities

$$
v = v_0 p_1(v), \quad \frac{v^2}{2} = \frac{v_0^2}{2}(\sqrt{2}p_2(v) + 1),
$$

or in terms of the matrix Q,

$$
Q = \begin{pmatrix} 1 & 0 & 0 \\ 0 & v_0 & 0 \\ \frac{v_0^2}{2} & 0 & \frac{v_0^2}{\sqrt{2}} \end{pmatrix}.
$$

The derivative coefficients, given in matrix form in (4.2.6), are

$$
d_{10} = \frac{1}{v_0}, \quad d_{20} = 0, \quad d_{21} = \frac{\sqrt{2}}{v_0}.
$$

We search for a solution $f = \mathcal{N} + g$ where g is expanded in terms of the first $M+1$ Hermite polynomials,

$$
g(x, v, t) = \sum_{n=0}^{M} w(v) He_n\left(\frac{v}{v_0}\right) g_n(x, t) = \boldsymbol{H}_M^T \boldsymbol{g}.
$$

We have expressed the sum in our preferred notation, which considers the sequence of asymmetrically weighted Hermite polynomials as a vector which may be combined via a dot product with the vector of coefficients, $\boldsymbol{g}(x, t)$.

The weighted inner product $\langle \cdot \rangle_{w^{-1}(v)}$ is discretized as the discrete dot product of two coefficient vectors, as demonstrated by the following sequence of identities:

$$\begin{aligned}\int_{\Omega_v} w^{-1}(v)g(v)h(v)\,\mathrm{d}v &= \int_{\Omega_v} w^{-1}(v)\left(\sum_{n=0}^{M} w(v)He_n\left(\frac{v}{v_0}\right)g_n\right)\left(\sum_{m=0}^{M} w(v)He_m\left(\frac{v}{v_0}\right)h_m\right)\mathrm{d}v \\ &= \sum_{n=0}^{M}\sum_{m=0}^{M} g_n h_m \int_{\Omega_v} w(v)He_n\left(\frac{v}{v_0}\right)He_m\left(\frac{v}{v_0}\right)\mathrm{d}v \\ &= \sum_{n=0}^{M}\sum_{m=0}^{M} g_n h_m \delta_{nm} \\ &= \boldsymbol{g}^T\boldsymbol{h}.\end{aligned}$$

As we have seen in (4.2.5) and (4.2.6), the multiplication-by-v and differentiation operators in the Hermite basis are easily represented by tridiagonal matrices. Finally, the projection P_Φ can be discretized as the operation which discards all but the first three Hermite coefficients, and conversely $P_\Phi^\perp$ as the operation which sets the first three Hermite coefficients to zero:

$$P_\Phi f(v) = \boldsymbol{H}_M^T\begin{pmatrix} f_0 \\ f_1 \\ f_2 \\ 0 \\ 0 \\ \vdots \end{pmatrix}, \quad P_\Phi^\perp f(v) = \boldsymbol{H}_M^T\begin{pmatrix} 0 \\ 0 \\ 0 \\ f_3 \\ f_4 \\ \vdots \end{pmatrix} \tag{8.5.1}$$

The discrete operators (4.2.5), (4.2.6), and (8.5.1) completely specify a velocity space discretization of our scheme.

In order to avoid numerical instability due to the Gibbs phenomenon, we apply a filter to the vector of Hermite modes after each timestep. The filtering procedure is the Hou-Li filter, described in Section 4.4.

8.5.2 Truncated domain finite difference method with Legendre weight

While global spectral methods for velocity space such as the Hermite method exhibit high accuracy and are easy to implement, they are not the only option available, nor the best in all circumstances. Grid-based methods, in particular Discontinuous Galerkin methods, are often preferred for their better resolution of fine filamentation structures in velocity space [59, 51]. In particular, unlike Hermite spectral methods, grid-based methods do not suffer from degraded resolution when the drift velocity or temperature of the local solution differs too much from the "reference" velocity and temperature around which the polynomial basis is expanded.

To illustrate the flexibility of our scheme to accomodate a variety of velocity discretizations including grid-based methods, in this section we describe a finite difference discretization of a truncated velocity space with the macro-micro decomposition based on the Legendre polynomials $P_i(v)$.

Our truncated velocity domain is $\Omega_v = [-v_{max}, v_{max}]$, with the constant weight function

$$w(v) = \frac{1}{v_{max}}.$$

The Legendre polynomials scaled to Ω_v satisfy the orthogonality relation

$$\int_{\Omega_v} w(v) P_n\left(\frac{v}{v_{max}}\right) P_m\left(\frac{v}{v_{max}}\right) \mathrm{d}v = \delta_{nm}.$$

The recurrence relation for the orthonormal Legendre polynomials has coefficients

$$a_n = \frac{n+1}{\sqrt{(2n+1)(2n+3)}} v_{max}, \quad b_n = 0,$$

and the first several examples are

$$\begin{aligned} p_0(v) &= P_0\left(\frac{v}{v_{max}}\right) = \sqrt{\frac{1}{2}} & p_2(v) &= P_2\left(\frac{v}{v_{max}}\right) = \sqrt{\frac{5}{8}}\left(3\left(\frac{v}{v_{max}}\right)^2 - 1\right) \\ p_1(v) &= P_1\left(\frac{v}{v_{max}}\right) = \sqrt{\frac{3}{2}}\frac{v}{v_{max}} & p_3(v) &= P_3\left(\frac{v}{v_{max}}\right) = \sqrt{\frac{7}{8}}\left(5\left(\frac{v}{v_{max}}\right)^3 - 3\frac{v}{v_{max}}\right). \end{aligned}$$

From the first three Legendre polynomials it is easy to verify the identities

$$v = \sqrt{\frac{2}{3}} v_{max} p_1(v), \quad \frac{v^2}{2} = \frac{v_{max}^2}{6}\left(\sqrt{\frac{8}{5}} p_2(v) + 1\right),$$

or in terms of the matrix Q,

$$Q = \begin{pmatrix} \sqrt{2} & 0 & 0 \\ 0 & \sqrt{\frac{2}{3}} v_{max} & 0 \\ \frac{v_{max}^2}{6} & 0 & \sqrt{\frac{8}{5}}\frac{v_{max}^2}{6} \end{pmatrix}.$$

The derivative coefficients d_{10}, d_{20}, d_{21} are as follows:

$$d_{10} = \frac{\sqrt{3}}{v_{max}}, \quad d_{20} = 0, \quad d_{21} = \frac{\sqrt{15}}{v_{max}}.$$

We discretize velocity space using a finite difference discretization. The hyperbolic term in the L step, equation (8.3.48), is discretized using a piecewise-linear MUSCL reconstruction [104] with the monotonized-central (MC) slope limiter [103] and a Lax-Friedrichs numerical flux [77, 78]. Second-order derivatives in v are discretized using a second-order centered finite difference operator. Inner products in v are computed using the midpoint rule. Boundary conditions on g are implemented via extrapolation of the Dirichlet boundary condition $g(x, v_b, t) = -\mathcal{N}(x, v, t)$ into a ghost cell layer that is two cells wide.

8.6 Numerical results

In this section we provide numerical results from standard benchmark problems in computational plasma physics. All benchmarks are implemented in both the global Hermite and finite difference velocity discretizations. For the Hermite discretization we use $v_0 = 1.0$ corresponding to a weight function $w(v) = \frac{1}{\sqrt{2\pi}} e^{-\frac{v^2}{2}}$. For the finite difference discretization, we use $v_{max} = 8.0$. Except where otherwise indicated, we use the second-order time integrator and neglect collisions, $\nu = 0.0$.

8.6.1 Verification of second-order temporal accuracy

To verify the claimed second-order accuracy of the splitting scheme described in Section 8.3.3, we perform a convergence study using the weak Landau damping numerical test of Section 8.6.2. This is solved using a fifth-order WENO finite difference discretization [92] in space with $N_x = 128$ grid points, and the Hermite spectral discretization in velocity with $M = 256$. The first-order scheme is run with Δt ranging from 4×10^{-3} to 1.25×10^{-4}, while the second-order scheme is run with Δt from 8×10^{-3} to 5×10^{-4}. Convergence is observed by comparing the solution $f(\Delta t)$ with the refined solution $f(\Delta t/2)$ at time $t = 5.0$, and taking the L^2 norm of the difference. The results are shown in Figure 8.1. We observe excellent agreement between the theoretical and empirical rates of convergence for both integrators.

8.6.2 Weak Landau damping

As a physics test, we reproduce the ubiquitous weak Landau damping benchmark problem with wavenumber $k = 0.5$ using both the Hermite spectral discretization and the Legendre-weighted finite difference discretization in velocity.

The initial condition is

$$f_0(x, v) = (1 + \delta \cos(kx)) e^{-v^2/2}, \quad x \in (0, 2\pi/k).$$

The perturbation size is set to $\delta = 1 \times 10^{-3}$. The domain is discretized with $N_x = 128$ points in the x direction and with either $M = 256$ Hermite modes or $N_v = 256$ velocity grid points. The rank is set to $r = 6$. We run the simulation with the second-order integrator to time $t = 40$, using timesteps of $\Delta t = 2 \times 10^{-3}$. The result is shown in Figure 8.2. We measure a damping rate of $\gamma = -0.1525$ for the Hermite spectral discretization

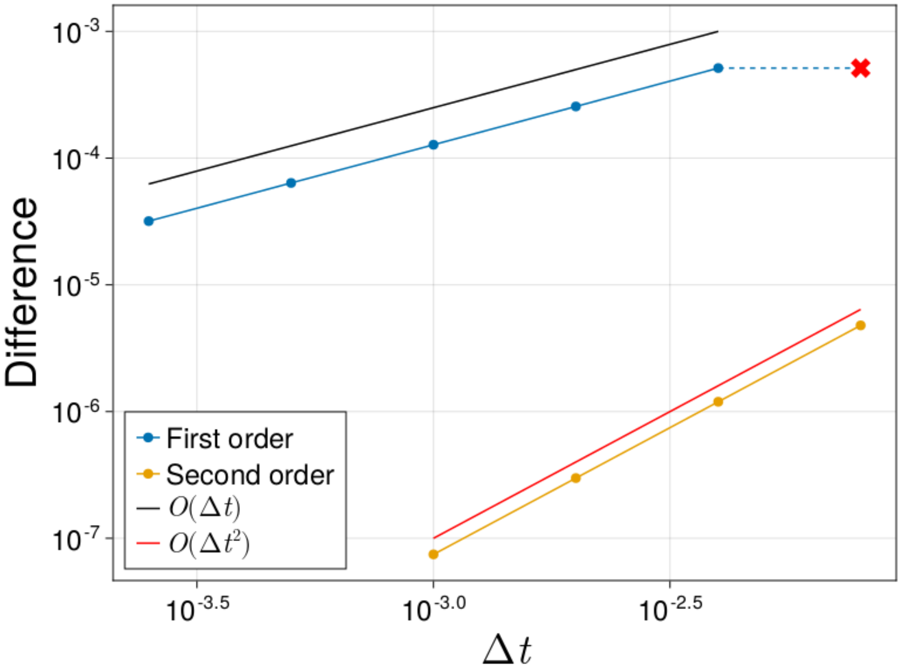

Figure 8.1: Convergence plot demonstrating first and second-order accuracy in time of the respective time integrators. The first-order integrator is unstable for $\Delta t = 8 \times 10^{-3}$.

and $\gamma = -0.1523$ for the finite difference discretization of velocity space, demonstrating good agreement with the linear theory prediction of $\gamma = -0.153$.

Collisional Landau damping

We validate the correctness of our solver including collisionality by comparing the Landau damping phenomenon at a variety of collision frequencies ν. We solve the same weak Landau damping problem as above, but with the collision frequency ν set to $0.0, 0.25$, and 1.0. The Hermite spectral solver is run with $M = 256$ and $\Delta t = 2 \times 10^{-3}$, the same as the collisionless example. In contrast, the finite difference velocity discretization becomes more stiff as the diffusive collision term grows larger, so for that discretization we reduce both the velocity grid spacing and timestep to $N_v = 128$ grid points and $\Delta t = 5 \times 10^4$. The results are shown in Figure 8.3.

8.6.3 Strong Landau damping

In this example we use the same wavenumber and domain as the weak Landau damping problem, $k = 0.5$, but set $\delta = 0.5$ to explore the strong (nonlinear) Landau damping regime. Again we present results from both the Hermite spectral and finite difference discretizations in velocity space. The simulation is run with $r = 16$, on a grid with $N_x = 128$ grid points in x, and either $M = 256$ Hermite modes or $N_v = 256$ velocity grid points. The timestep is set to $\Delta t = 4 \times 10^{-3}$, and the initial condition is evolved using the second-order

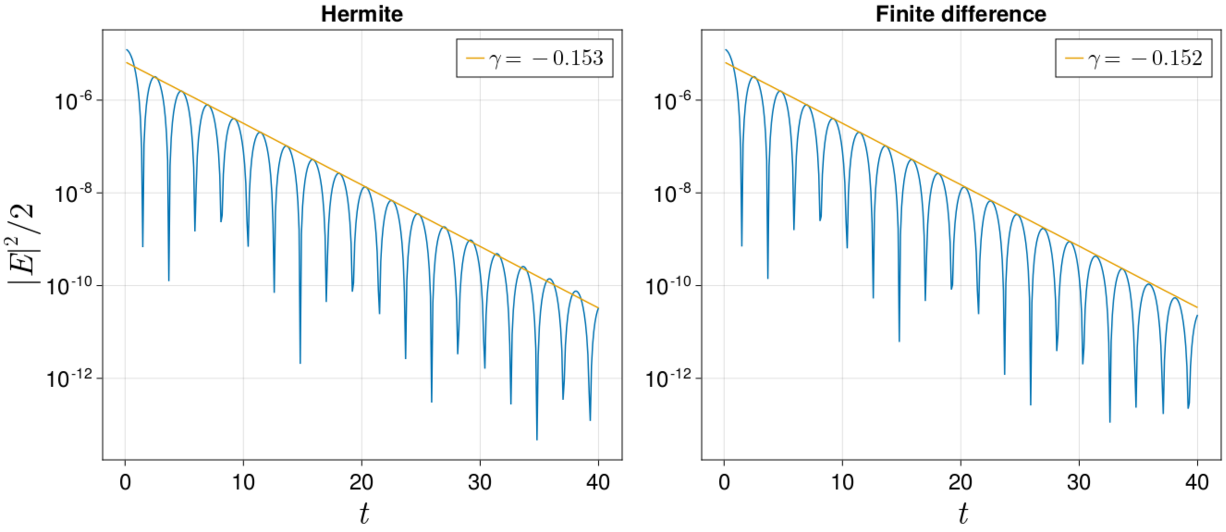

(a) Electric energy and best linear fit.

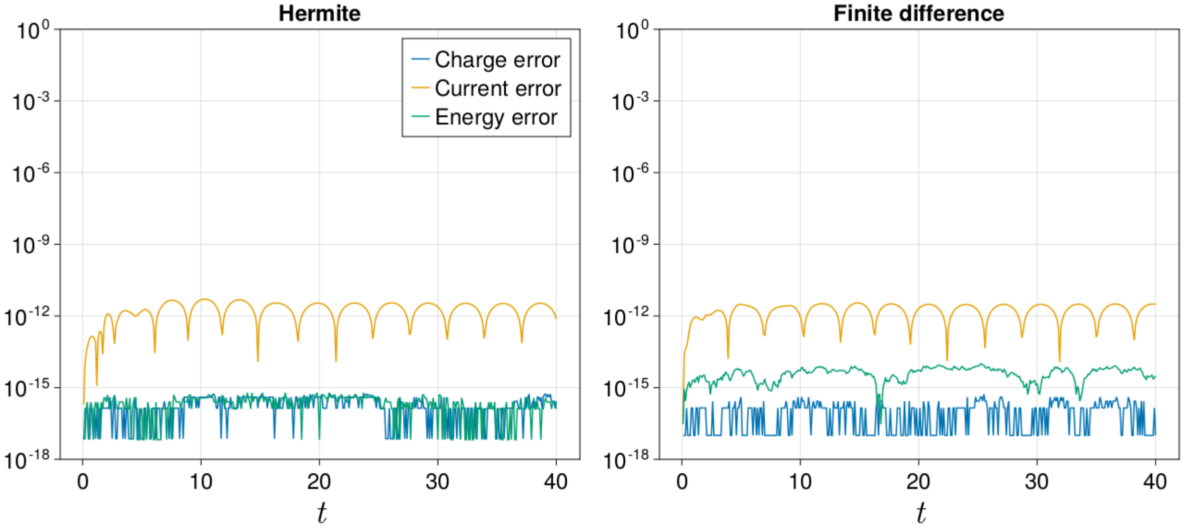

(b) Charge, current and energy conservation error.

Figure 8.2: Weak Landau damping example, demonstrating exact conservation of charge and energy with the second-order time integrator.

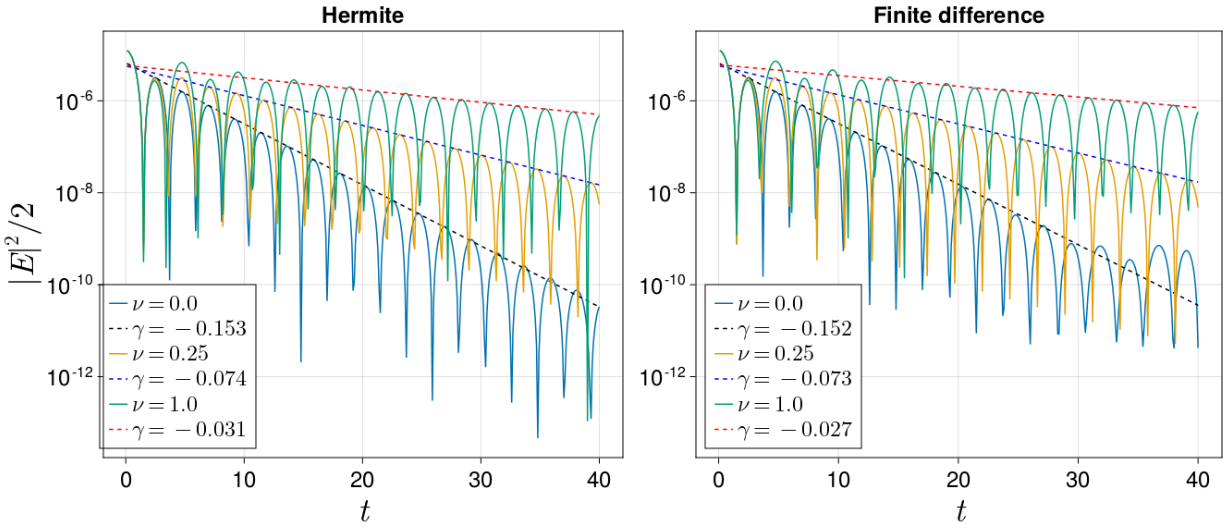

(a) Measured damping rates as a function of collision frequency ν. Increasing collisionality brings the solution nearer to the fluid limit, which experiences no damping. Our results agree well with published results in [52].

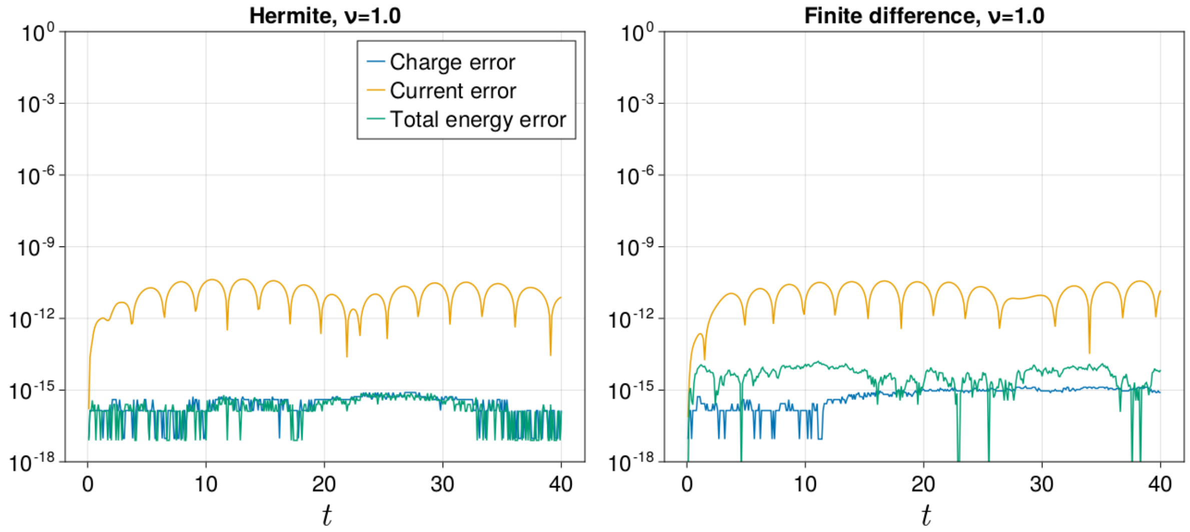

(b) Conservation of charge and energy by the second-order integrator, tested in the $\nu = 1.0$ case.

Figure 8.3: Collisional Landau damping example.

energy-conserving integrator to $t = 50.0$. The results are shown in Figure 8.4, including the phase space density at $t = 25.0$. Conservation properties of both first and second order integrators on this strong Landau damping problem are shown in Figure 8.5.

8.6.4 Two-stream instability

Here we reproduce the two-stream instability example from [39]:

$$f(x, v, 0) = \frac{2}{7}(1 + 5v^2)(1 + \delta((\cos(2kx) + \cos(3kx))/1.2 + \cos(kx)))\frac{1}{\sqrt{2\pi}}e^{-v^2/2},$$

with $\delta = 0.01, k = 0.5$. This form of the distribution function is chosen to give the following analytic forms for the zeroth and second Hermite moments:

$$\begin{aligned} f_0^{\text{Hermite}}(x, 0) &= \frac{12}{7}(1 + \delta((\cos(2kx) + \cos(3kx))/1.2 + \cos(kx))), \\ f_2^{\text{Hermite}}(x, 0) &= \frac{10\sqrt{2}}{7}(1 + \delta((\cos(2kx) + \cos(3kx))/1.2 + \cos(kx))). \end{aligned}$$

We run this simulation on a grid with $N_x = 256$ grid points in x. For the Hermite spectral discretization we use $M = 256$ Hermite modes, and for the finite difference discretization we use $N_v = 256$ velocity grid points. The rank is set to $r = 20$, and the instability is evolved with $\Delta t = 4 \times 10^{-3}$ well into the nonlinear phase, up to $t = 50.0$. The results are shown in Figure 8.6.

8.7 Extensions to the single-species 1D1V scheme

8.7.1 Multispecies conservation

In multispecies plasmas, interspecies collisions mean that conservation laws are no longer satisifed by individual species distribution functions, but rather the whole plasma system. That is, the single-species conservation laws (8.1.6), (8.1.7), and (8.1.8) are replaced by the following equations:

$$\partial_t \rho_a + \nabla \cdot (\rho_a \boldsymbol{u}_a) = 0, \tag{8.7.1}$$

$$\sum_a (\partial_t(\rho_a \boldsymbol{u}_a) + \nabla \cdot \sigma_a) = \sum \rho_a \boldsymbol{E}, \tag{8.7.2}$$

$$\sum_a (\partial_t \kappa_a + \nabla \mathbf{q}_a) = \sum_a \frac{Z_a}{A_a} \rho_a \boldsymbol{u}_a \cdot \boldsymbol{E}, \tag{8.7.3}$$

where

$$\rho_a = A_a \int_{\mathbb{R}^d} f_a \, \mathrm{d}v, \quad \kappa_a = A_a \int_{\mathbb{R}^d} \frac{|\boldsymbol{v}|^2}{2} f_a \, \mathrm{d}v,$$

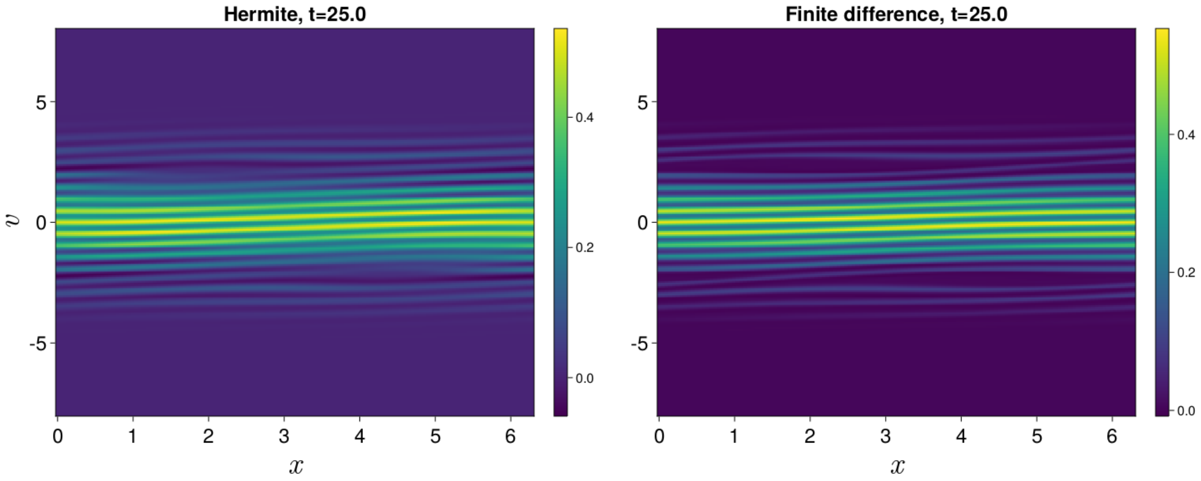

(a) Phase space distribution $f(x, v, t = 25.0)$ showing strong filamentation. The finite difference discretization is both less diffuse and shows less severe positivity violations than the Hermite discretization, which demonstrates the importance of compatibility with a diversity of velocity space discretizations.

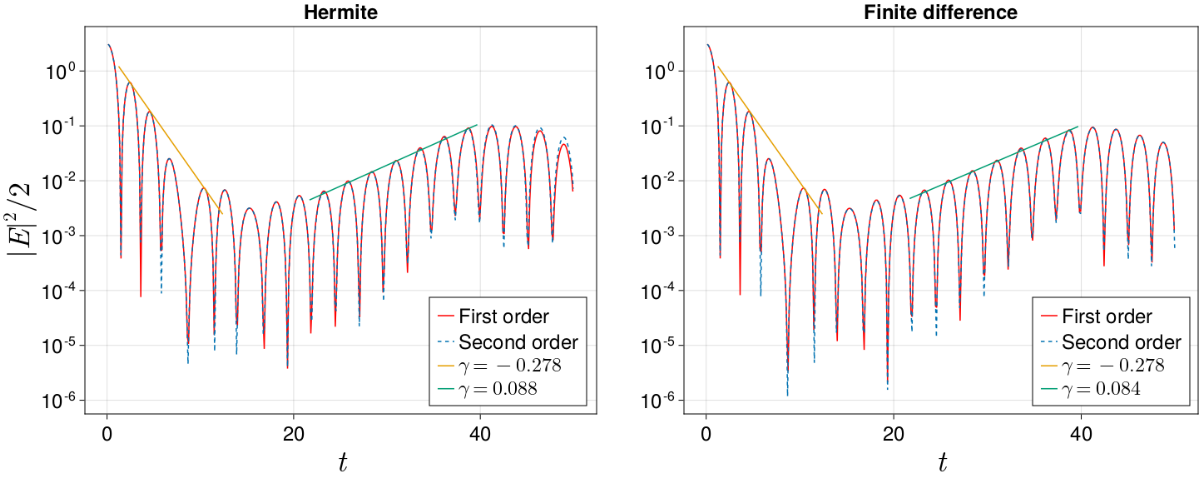

(b) Electric energy traces. Our implementation matches both the initial nonlinear damping and subsequent growth rates, compared to published results such as [59].

Figure 8.4: Strong Landau damping example.

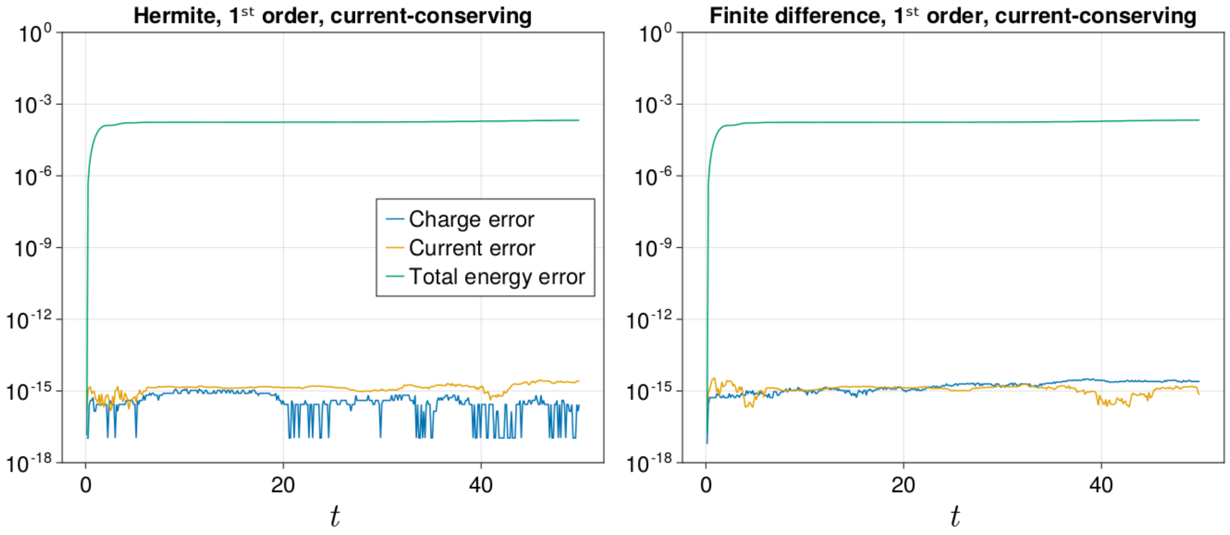

(a) Exact conservation of charge and current by the first order integrator with a Poisson solve.

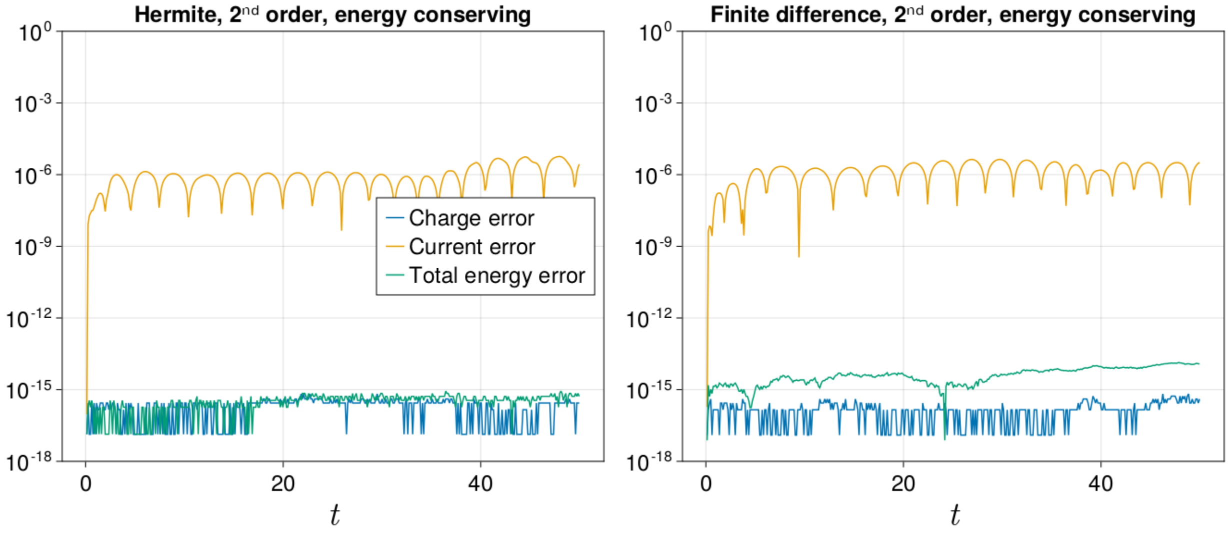

(b) Exact conservation of charge and energy by the second order integrator with an Ampère solve.

Figure 8.5: Conservation plots in the strong Landau damping example. We observe the benefit of overall second-order accuracy in the improved conservation of current in (b), compared to the conservation error of energy in (a).

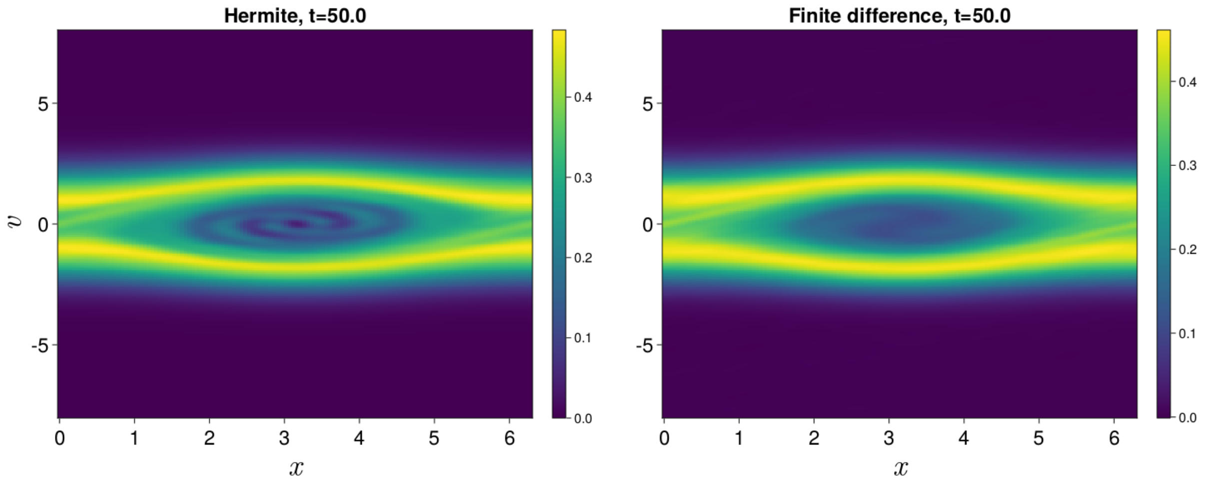

(a) Plots of $f(x, v, t = 50.0)$, depicting nonlinear phase space trapping in the two-stream instability.

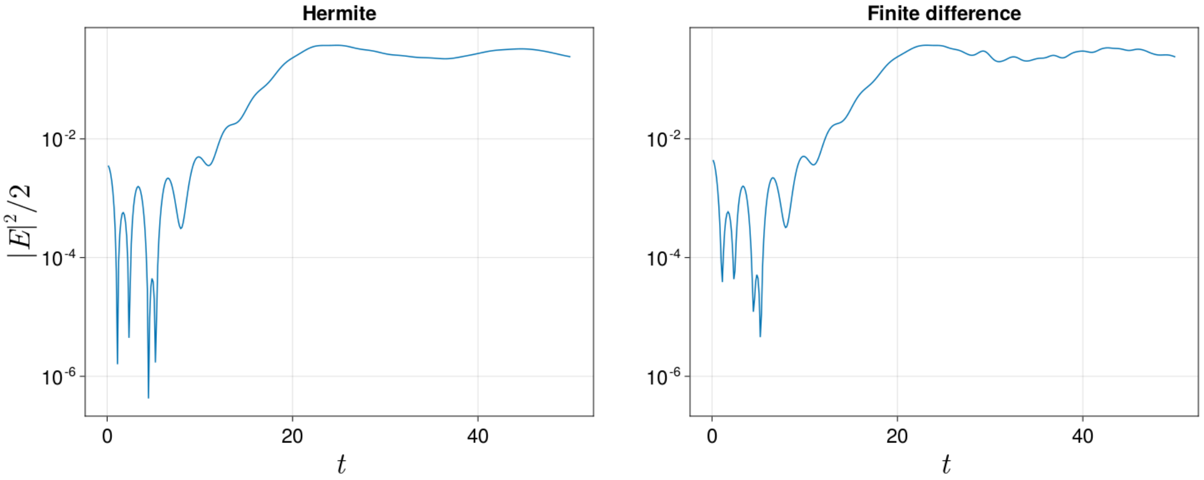

(b) Electric energy trace. Our results show excellent agreement with published numerical results in [39].

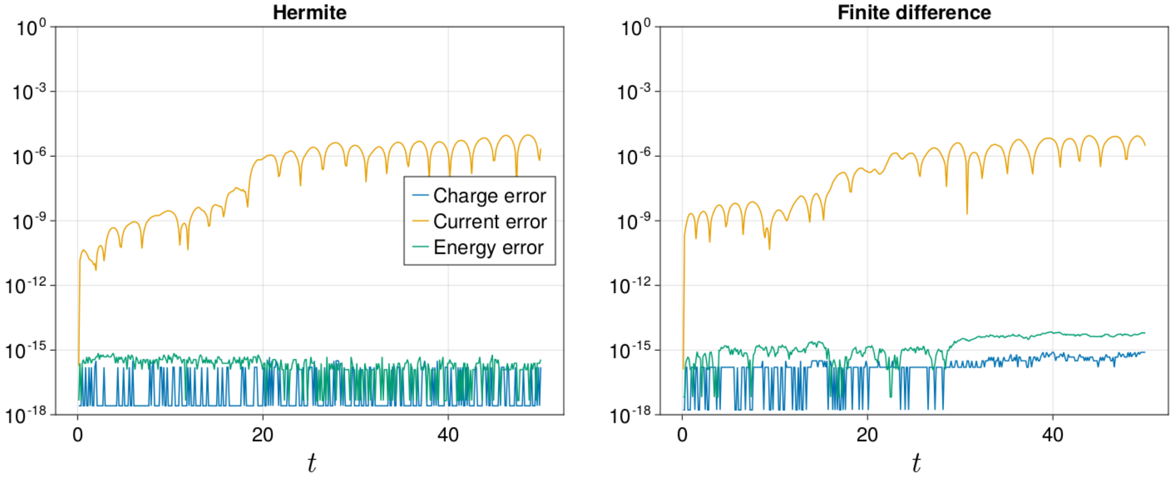

(c) Exact conservation of charge and energy by the second order integrator.

Figure 8.6: Two-stream instability.

and

$$\sigma_a = A_a \int_{\mathbb{R}^d} (\boldsymbol{v} \otimes \boldsymbol{v}) f_a \, \mathrm{d}v, \quad \mathbf{q}_a = \frac{A_a}{2} \int_{\mathbb{R}^d} \boldsymbol{v} |\boldsymbol{v}|^2 f_a \, \mathrm{d}v.$$

Despite the additional complexity, the macro-micro decomposition described here can accomodate multi-species conservation with minimal modification. The multi-species conservation statements (8.7.1), (8.7.2), and (8.7.3) can be reformulated as statements about the linear combination of the orthogonal moments f_0, f_1, f_2 for each species a via the existence of a mapping matrix Q. As in the single-species case, conservation properties are therefore statements about the macroscopic coefficients, in which the DLR discretization of the microscopic equation plays no role. Local conservation can then be achieved by suitable discretization of the macro equations. We leave the detailed derivation of such a scheme to future work.

8.7.2 Multi-dimensional conservative DLR

The extension of the method described in this chapter to multiple dimensions is fairly straightforward, since all discretizations are based on one-dimensional primitives that fit naturally into a tensor-product discretization of 2 or 3 physical and velocity space dimensions. The lone subtlety relates to the definition of the projections P_Φ and $P_\Phi^\perp$. In particular, the particle energy in multiple dimensions is defined as the moment of a linear combination of one-dimensional polynomials:

$$\kappa = \int \frac{|\boldsymbol{v}|^2}{2} f \, \mathrm{d}v = \int \frac{v_x^2 + v_y^2 + v_z^2}{2} f \, \mathrm{d}v.$$

Therefore, the simple truncation projection described in (8.5.1) for the Hermite basis will not work in multiple dimensions. In this section we briefly sketch the solution for the Hermite expansion in 2D2V.

The Hermite spectral discretization described in Chapter 4 results in an array of unknowns which are the Hermite polynomial moments of f. In 2D2V, these unknowns can be arranged into a vector in order of total polynomial degree, so that

$$\boldsymbol{f} = (f_{00}, f_{10}, f_{01}, f_{20}, f_{02}, f_{11}, \dots)^T.$$

The continuous solution f is then

$$f(\boldsymbol{x}, \boldsymbol{v}, t) = w(\boldsymbol{v}) \boldsymbol{p}(\boldsymbol{v})^T \boldsymbol{f}(\boldsymbol{x}, t),$$

where we have defined the vector of Hermite polynomials

$$\boldsymbol{p}(\boldsymbol{v}) = \begin{pmatrix} p_{00}(\boldsymbol{v}) \\ p_{10}(\boldsymbol{v}) \\ p_{01}(\boldsymbol{v}) \\ p_{20}(\boldsymbol{v}) \\ \vdots \end{pmatrix} = \begin{pmatrix} He_0(v_x/v_0)He_0(v_z/v_0) \\ He_1(v_x/v_0)He_0(v_z/v_0) \\ He_0(v_x/v_0)He_1(v_z/v_0) \\ He_2(v_x/v_0)He_0(v_z/v_0) \\ \vdots \end{pmatrix}$$

and the weight function

$$w(\boldsymbol{v}) = \frac{1}{2\pi v_0^2} e^{-\frac{|\boldsymbol{v}|^2}{2v_0^2}}.$$

The macroscopic part $\mathcal{N}$ must have the following form,

$$\mathcal{N} = w(v)[p_{00}(\boldsymbol{v})f_{00} + p_{10}(\boldsymbol{v})f_{10} + p_{01}(\boldsymbol{v})f_{01} + p_2(\boldsymbol{v})f_2],$$

where the polynomial p_2 is associated with the total energy, and is defined

$$p_2 = \frac{p_{20}(\boldsymbol{v}) + p_{02}(\boldsymbol{v})}{2},$$

and the combined second moment is $f_2 = f_{20} + f_{02}$. Thus, the action of the projection operator P_Φ on f can be written in the following matrix form:

$$\mathcal{N} = P_\Phi f = w(\boldsymbol{v})\boldsymbol{p}(\boldsymbol{v})^T \begin{pmatrix} 1 & & & & & & \\ & 1 & & & & & \\ & & 1 & & & & \\ & & & 1/2 & 1/2 & & \\ & & & 1/2 & 1/2 & & \\ & & & & & 0 & \\ & & & & & & \ddots \end{pmatrix} \boldsymbol{f}. \tag{8.7.4}$$

The microscopic part g is the image of the orthogonal complement of P_Φ,

$$g = P_\Phi^\perp f = w(\boldsymbol{v})\boldsymbol{p}(\boldsymbol{v})^T \begin{pmatrix} 0 & & & & & & \\ & 0 & & & & & \\ & & 0 & & & & \\ & & & 1/2 & -1/2 & & \\ & & & -1/2 & 1/2 & & \\ & & & & & 1 & \\ & & & & & & \ddots \end{pmatrix} \boldsymbol{f}. \tag{8.7.5}$$

To summarize, in two dimensions we define

$$\mathcal{N} = w(v)[p_0(\boldsymbol{v})f_0 + p_{10}(\boldsymbol{v})f_{10} + p_{01}(\boldsymbol{v})f_{01} + p_2(\boldsymbol{v})f_2],$$

where

$$p_2 = \frac{p_{20}(\boldsymbol{v}) + p_{02}(\boldsymbol{v})}{2},$$

and $f_2 = \langle p_2 f \rangle_v$. It can be seen that this definition is independent of the choice of polynomial family p and inner product; thus, the same formulation works for the truncated domain Legendre finite difference discretization as long as the appropriate inner product is used to calculate the moments f_0, f_{10}, f_{01}, f_2. The extension to three dimensions is simple.

Chapter 9

Conclusion

This thesis has been organized around a pair of interlocking themes. The first is the familiar dilemma faced by computational physicists and applied mathematicians: affordable physical models are inaccurate, but the accurate models are too costly for the available numerical methods. This dilemma organizes our research into reduced models for plasma physics in two directions. By seeking physically accurate extensions to well-known reduced models, we are able to extend their regimes of validity and incorporate important micro-scale physics in an affordable way. This is the direction of the work presented in Chapters 3-5. On the other hand, the computational physicist's dilemma can be resolved from the other direction: if kinetic models may be solved quickly and accurately enough, there is no need to bother with deriving reduced models. This research horizon has driven our engagement with and contributions to the literature on dynamical low-rank methods, which were the focus of Chapters 6-8. We use the word "horizon" with consideration: there will never be a numerical method for kinetics so excellent that there is no need for fluid models, but it is a good direction to strive towards.

Despite the power of asymptotic methods to determine leading-order effects, our work in this thesis has underlined the crucial lesson that practitioners should exercise caution when applying them. Translating the complex nonlinear dynamics of a plasma into a single value for ϵ is highly nontrivial. And even when the plasma dynamics exhibit an ordering that make an asymptotic theory formally valid, it may be the case that the theory does not take advantage of all of the "smallness" available. The disagreement between the drift ordering closure and Braginskii closure, explained in Chapter 3 and explored numerically in Chapter 5, is an example of this. By taking advantage of the smallness of the fluid velocity in a way that the Braginskii closure does not, we found that the leading-order gyroviscous stress is influenced by both velocity-driven shear stress and heat flux-driven distortions to the distribution function. The Braginskii asymptotic theory therefore breaks down not only when inertial and electrostatic forces are strong compared to magnetic forces, but also when these forces are weak compared to magnetic forces, in which case the leading-order linear relationship can be replaced by a more accurate model.

The advent of high-dimensional continuum kinetic simulation provides a powerful check for asymptotic methods. By running highly accurate kinetic simulations, we obtain a ground truth for the true shape of a kinetic distribution function, providing data on the magnitude and direction of kinetic closure terms. These methods enable evaluation of the accuracy and reliability of asymptotic transport closures, as well as an idea of the nature of higher order closure error terms. In Chapter 5 we performed such an investigation, which provided additional insight into the nature of the drift ordering closure for the magnetized Vlasov equation, including an affordable approximation which can be incorporated into existing codes. More broadly, continuum kinetic simulation methods can improve our understanding of kinetic effects and how they can be approximated by gathering data on a range of prototypical plasma features, such as sheared flows, gradients, $E \times B$ drifts, and more. This improved understanding feeds back into improved fluid models, which can be applied to solve problems which exhibit instances of the prototypical features that the closure has been tested on.

The second theme of this thesis is contained in its title, which contrasts asymptotic model reduction methods to non-asymptotic methods. The ascendancy of asymptotic analysis over the last 70 years in applied mathematics and computational physics is only beginning to wane in importance compared to the research enabled by computer systems. Today's supercomputing clusters have capabilities which in the 1950's would have been unfathomable to the founders of the field of nuclear fusion energy, optimists who certainly regarded controlled nuclear fusion as an easier problem to solve than exascale computing. In the event, the multiscale challenges of plasma physics have not yielded so easily to asymptotic methods. Multiscale phenomena, particularly turbulence in all its forms, cross scale boundaries in such a way that it is very hard to rely on the scale separation inherent to asymptotic analysis.

It is interesting to compare the recent experience of another scientific field, artificial intelligence. A compelling lens through which to understand that field's recent strides is offered by Rich Sutton's essay *The Bitter Lesson* [97]. Sutton argues that time and again, AI researchers have resisted the apparent fact that human cleverness is not as important to progress as simply throwing more computational resources at the problem. Indeed, the most advanced AI systems as of 2024 have almost no "built-in smarts", consisting of fairly simple architectures trained on terabytes of data and millions of GPU-hours. Is there a bitter lesson for plasma physics as well? Will the great computational plasma physics breakthroughs of the next ten years come from ever-larger simulations run on millions of GPUs, or will they come from clever interpretation and physics analysis of smaller-scale simulations?

There is good reason to be skeptical that whole-device direct numerical simulation of the Vlasov-Boltzmann equation will lead to nuclear fusion breakthroughs. The fastest timescales in a plasma are simply far too fast for such an approach to be feasible. The straightforward dynamical low-rank methods described in Chapter 6 are a different sort of approach: they seek to overcome the curse of dimensionality by exploiting mathematical structures which are foreign to classical plasma physics, rather than blow past it with overwhelming computational resources. But the straightforward DLR methods are naïve of plasma physics concepts as simple as conservation of energy, let alone the more complex concepts that structure asymptotics-informed analyses. It might be a bitter pill if such a method were a one-size-fits-all solution to multiscale plasma modeling.

On the contrary, we believe that the advent of powerful non-asymptotic, physics-naïve model reduction techniques such as dynamical low-rank approximation increase rather than diminish the importance of careful asymptotic analysis. In our opinion the most fruitful avenue for continued development of DLR methods in kinetic plasma physics lies in viewing DLR and asymptotic, physics-informed methods as complementary. The thrust of the work described in Chapters 7 and 8 is to incorporate a pair of essential plasma physics concepts into DLR methods. In Chapter 7 we described a DLR method for a collisional plasma kinetic equation which smoothly approaches the fluid limit. The result is a method which, we anticipate, can be more easily integrated into a rank-adaptive system that can efficiently solve both fluid and kinetic dynamics. In Chapter 8, we described how to perform a macro-micro decomposition of the kinetic equation which is amenable to deriving the DLR equations of motion for the microscopic part. The macro-micro decomposition ensures conservation while the DLR approximation provides efficient representation of kinetic physics.

We have presented these contributions as modifications to a base DLR scheme. But it may be just as valid to view the algorithms in Chapters 7 and 8 as numerical methods for physics-informed reformulations of the kinetic equation, which use a DLR scheme as just one component. For example, the system of equations for the leading three moments in our conservative macro-micro decomposition, Equation (8.2.13), is equivalent to the classic five-moment system of equations in one dimension. From this point of view the role of the dynamical low-rank approximation is to provide a moment closure for the unclosed moment f^3. It is possible, therefore, to interpret this work as another entry in the long history of attempts to address the moment closure problem for fluid systems.

When considered from this point of view, numerous extensions to the conservative macro-micro decomposition immediately present themselves. We have stopped at three moments in one dimension because our aim was to achieve a conservative method. But the moment closure problem occurs at every possible stopping point in the moment hierarchy. Therefore, the same technique can perhaps be applied to obtain a DLR closure for a 10-moment or 13-moment model in three dimensions, providing closures for third- and fourth-order tensor moments of the distribution function.

Another asymptotic limit where DLR algorithms may be fruitfully applied is the strongly magnetized, collisionless limit explored in Chapters 3-5. As we saw in Chapter 3, in this regime the perpendicular and parallel velocity dependence can be separated through asymptotic arguments. Integrating over $\boldsymbol{v}_\perp$ to form the system of semi-fluid equations reduces the number of independent dimensions from six to four. A DLR algorithm applied to the parallel velocity dependence might greatly increase the performance of a numerical solver for these equations, while still capturing important kinetic physics in the parallel direction such as microinstabilities. Similarly, the nearly gyrotropic nature of strongly magnetized distribution functions means that the two components of perpendicular velocity can be separated from each other, since both $|\boldsymbol{v}_\perp|^2$ and $\exp(|\boldsymbol{v}_\perp|^2)$ are separable functions of the two perpendicular velocity coordinates. Thus, in addition to the more traditional asymptotic expansions enabled by the strongly magnetized regime, the structure that magnetization imposes on the distribution function may also make DLR approximations more promising.

The work presented in this thesis is only a very small portion of the research that remains to arrive at a satisfactory, comprehensive approach for solving problems in multiscale plasma dynamics. In addition to dynamical low-rank, other non-asymptotic approaches such as sparse grid and step-then-truncate low-rank methods are likely to be of great importance. On the asymptotic side, it is clear that collisionless drift ordering explored here is only one of many limits of physical interest, and that we have only scratched the surface in this work. Nevertheless, we hope that the techniques we have demonstrated and the point of view we have described are of some value to future researchers.

Bibliography

[1] H. Abu-Shawareb et al. "Lawson Criterion for Ignition Exceeded in an Inertial Fusion Experiment". In: *Physical Review Letters* 129.7 (Aug. 2022), p. 075001. ISSN: 0031-9007, 1079-7114. DOI: `10.1103/PhysRevLett.129.075001`. (Visited on 05/09/2024).

[2] Markus Bachmayr et al. "Existence of Dynamical Low-Rank Approximations to Parabolic Problems". In: *Mathematics of Computation* 90.330 (Mar. 2021), pp. 1799–1830. ISSN: 0025-5718, 1088-6842. DOI: `10.1090/mcom/3626`. (Visited on 05/14/2024).

[3] E. V. Belova. "Nonlinear Gyroviscous Force in a Collisionless Plasma". In: *Physics of Plasmas* 8.9 (Sept. 2001), pp. 3936–3944. ISSN: 1070-664X, 1089-7674. DOI: `10.1063/1.1389093`. (Visited on 05/14/2024).

[4] Mounir Bennoune, Mohammed Lemou, and Luc Mieussens. "Uniformly Stable Numerical Schemes for the Boltzmann Equation Preserving the Compressible Navier–Stokes Asymptotics". In: *Journal of Computational Physics* 227.8 (Apr. 2008), pp. 3781–3803. ISSN: 00219991. DOI: `10.1016/j.jcp.2007.11.032`. (Visited on 05/05/2023).

[5] Tim Besard, Christophe Foket, and Bjorn De Sutter. "Effective Extensible Programming: Unleashing Julia on GPUs". In: *IEEE Transactions on Parallel and Distributed Systems* 30.4 (Apr. 2019), pp. 827–841. ISSN: 1045-9219, 1558-2183, 2161-9883. DOI: `10.1109/TPDS.2018.2872064`. (Visited on 05/14/2024).

[6] Tim Besard et al. "Rapid Software Prototyping for Heterogeneous and Distributed Platforms". In: *Advances in Engineering Software* 132 (June 2019), pp. 29–46. ISSN: 09659978. DOI: `10.1016/j.advengsoft.2019.02.002`. (Visited on 05/14/2024).

[7] Jeff Bezanson et al. "Julia: A Fresh Approach to Numerical Computing". In: *SIAM Review* 59.1 (Jan. 2017), pp. 65–98. ISSN: 0036-1445, 1095-7200. DOI: `10.1137/141000671`. (Visited on 05/14/2024).

[8] P. L. Bhatnagar, E. P. Gross, and M. Krook. "A Model for Collision Processes in Gases. I. Small Amplitude Processes in Charged and Neutral One-Component Systems". In: *Physical Review* 94.3 (May 1954), pp. 511–525. ISSN: 0031-899X. DOI: `10.1103/PhysRev.94.511`. (Visited on 09/08/2022).

[9] Charles K. Birdsall and A. Bruce Langdon. *Plasma Physics via Computer Simulation.* New York: McGraw-Hill, 1985. ISBN: 978-0-07-005371-7.

[10] John P. Boyd. *Chebyshev and Fourier Spectral Methods.* 2nd ed., rev. Mineola, N.Y: Dover Publications, 2001. ISBN: 978-0-486-41183-5.

[11] S.I. Braginskii. "Transport Processes in a Plasma". In: *Reviews of Plasma Physics.* Vol. 1. 1965.

[12] A. J. Brizard and T. S. Hahm. "Foundations of Nonlinear Gyrokinetic Theory". In: *Reviews of Modern Physics* 79.2 (Apr. 2007), pp. 421–468. ISSN: 0034-6861, 1539-0756. DOI: `10.1103/RevModPhys.79.421`. (Visited on 05/28/2024).

[13] José Canosa, Jenö Gazdag, and J.E. Fromm. "The Recurrence of the Initial State in the Numerical Solution of the Vlasov Equation". In: *Journal of Computational Physics* 15.1 (May 1974), pp. 34–45. ISSN: 00219991. DOI: `10.1016/0021-9991(74)90067-9`. (Visited on 09/21/2022).

[14] Carlo Cercignani. *The Boltzmann Equation and Its Applications.* Ed. by F. John, J. E. Marsden, and L. Sirovich. Vol. 67. Applied Mathematical Sciences. New York, NY: Springer New York, 1988. ISBN: 978-1-4612-6995-3 978-1-4612-1039-9. DOI: `10.1007/978-1-4612-1039-9`. (Visited on 03/11/2021).

[15] Carlo Cercignani et al. "Device Benchmark Comparisons via Kinetic, Hydrodynamic, and High-Hield Models". In: *Computer Methods in Applied Mechanics and Engineering* 181.4 (Jan. 2000), pp. 381–392. ISSN: 00457825. DOI: `10.1016/S0045-7825(99)00186-3`. (Visited on 08/30/2022).

[16] Gianluca Ceruti and Christian Lubich. "An Unconventional Robust Integrator for Dynamical Low-Rank Approximation". In: *BIT Numerical Mathematics* 62.1 (Mar. 2022), pp. 23–44. ISSN: 0006-3835, 1572-9125. DOI: `10.1007/s10543-021-00873-0`. (Visited on 08/30/2022).

[17] G.F. Chew, M.L. Goldberger, and F.E. Low. "The Boltzmann Equation and the One-Fluid Hydromagnetic Equations in the Absence of Particle Collisions". In: *Proceedings of the Royal Society of London.* A 236.1204 (July 1956), pp. 112–118.

[18] Bernardo Cockburn and Chi-Wang Shu. "Runge–Kutta Discontinuous Galerkin Methods for Convection-Dominated Problems". In: *Journal of Scientific Computing* 16.3 (2001), pp. 173–261. ISSN: 08857474. DOI: `10.1023/A:1012873910884`. (Visited on 05/14/2024).

[19] Jack Coughlin and Jingwei Hu. "Efficient Dynamical Low-Rank Approximation for the Vlasov-Ampère-Fokker-Planck System". In: *Journal of Computational Physics* (Sept. 2022), p. 111590. ISSN: 00219991. DOI: `10.1016/j.jcp.2022.111590`. (Visited on 09/09/2022).

[20] Jack Coughlin, Jingwei Hu, and Uri Shumlak. "Robust and Conservative Dynamical Low-Rank Methods for the Vlasov Equation via a Novel Macro-Micro Decomposition". In: *Journal of Computational Physics* 509 (July 2024), p. 113055. ISSN: 00219991. DOI: `10.1016/j.jcp.2024.113055`. (Visited on 05/08/2024).

[21] D. W. Crews and U. Shumlak. *Phase Space Eigenfunctions with Applications to Continuum Kinetic Simulations.* Feb. 2024. arXiv: `2402.02180 [physics]`. (Visited on 05/05/2024).

[22] Daniel W Crews. "Numerical Simulation of Collisionless Kinetic Plasma Turbulence". PhD thesis. University of Washington, 2022.

[23] I. A. M. Datta, D. W. Crews, and U. Shumlak. “Electromagnetic Extension of the Dory–Guest–Harris Instability as a Benchmark for Vlasov–Maxwell Continuum Kinetic Simulations of Magnetized Plasmas”. In: *Physics of Plasmas* 28.7 (July 2021), p. 072112. ISSN: 1070-664X, 1089-7674. DOI: `10.1063/5.0057230`. (Visited on 05/05/2024).

[24] I.A.M. Datta and U. Shumlak. “Computationally Efficient High-Fidelity Plasma Simulations by Coupling Multi-Species Kinetic and Multi-Fluid Models on Decomposed Domains”. In: *Journal of Computational Physics* 483 (June 2023), p. 112073. ISSN: 00219991. DOI: `10.1016/j.jcp.2023.112073`. (Visited on 01/17/2024).

[25] Pierre Degond. “Asymptotic Continuum Models for Plasmas and Disparate Mass Gaseous Binary Mixtures”. In: *Material Substructures in Complex Bodies*. Elsevier, 2007, pp. 1–62. ISBN: 978-0-08-044535-9. DOI: `10.1016/B978-008044535-9/50002-9`. (Visited on 10/28/2022).

[26] G.L. Delzanno. “Multi-Dimensional, Fully-Implicit, Spectral Method for the Vlasov–Maxwell Equations with Exact Conservation Laws in Discrete Form”. In: *Journal of Computational Physics* 301 (Nov. 2015), pp. 338–356. ISSN: 00219991. DOI: `10.1016/j.jcp.2015.07.028`. (Visited on 01/29/2021).

[27] F. W. J. Olver et al. *NIST Digital Library of Mathematical Functions*. 2020.

[28] William Douglass Dorland. “GYROFLUID MODELS OF PLASMA TURBULENCE”. PhD thesis. Princeton University, Nov. 1993.

[29] R. A. Dory, G. E. Guest, and E. G. Harris. “Unstable Electrostatic Plasma Waves Propagating Perpendicular to a Magnetic Field”. In: *Physical Review Letters* 14.5 (Feb. 1965), pp. 131–133. ISSN: 0031-9007. DOI: `10.1103/PhysRevLett.14.131`. (Visited on 06/03/2021).

[30] J. P. Dougherty. “Model Fokker-Planck Equation for a Plasma and Its Solution”. In: *Physics of Fluids* 7.11 (1964), p. 1788. ISSN: 00319171. DOI: `10.1063/1.2746779`. (Visited on 04/07/2022).

[31] Dale R. Durran. *Numerical Methods for Fluid Dynamics*. Vol. 32. Texts in Applied Mathematics. New York, NY: Springer New York, 2010. ISBN: 978-1-4419-6411-3 978-1-4419-6412-0. DOI: `10.1007/978-1-4419-6412-0`. (Visited on 02/08/2024).

[32] Lukas Einkemmer. “Accelerating the Simulation of Kinetic Shear Alfvén Waves with a Dynamical Low-Rank Approximation”. In: (2023). DOI: `10.48550/ARXIV.2306.17526`. (Visited on 09/27/2023).

[33] Lukas Einkemmer, Jingwei Hu, and Lexing Ying. “An Efficient Dynamical Low-Rank Algorithm for the Boltzmann-BGK Equation Close to the Compressible Viscous Flow Regime”. In: *SIAM Journal on Scientific Computing* 43.5 (Jan. 2021), B1057–B1080. ISSN: 1064-8275, 1095-7197. DOI: `10.1137/21M1392772`. (Visited on 04/13/2022).

[34] Lukas Einkemmer and Ilon Joseph. “A Mass, Momentum, and Energy Conservative Dynamical Low-Rank Scheme for the Vlasov Equation”. In: *Journal of Computational Physics* 443 (Oct. 2021), p. 110495. ISSN: 00219991. DOI: `10.1016/j.jcp.2021.110495`. (Visited on 04/13/2022).

[35] Lukas Einkemmer, Jonas Kusch, and Steffen Schotthöfer. "Conservation Properties of the Augmented Basis Update & Galerkin Integrator for Kinetic Problems". In: (Nov. 2023). DOI: `10.48550/arXiv.2311.06399`. arXiv: `2311.06399 [cs, math]`. (Visited on 02/20/2024).

[36] Lukas Einkemmer and Christian Lubich. "A Low-Rank Projector-Splitting Integrator for the Vlasov–Poisson Equation". In: *SIAM Journal on Scientific Computing* 40.5 (Jan. 2018), B1330–B1360. ISSN: 1064-8275, 1095-7197. DOI: `10.1137/18M116383X`. (Visited on 03/11/2022).

[37] Lukas Einkemmer and Christian Lubich. "A Quasi-Conservative Dynamical Low-Rank Algorithm for the Vlasov Equation". In: *SIAM Journal on Scientific Computing* 41.5 (Jan. 2019), B1061–B1081. ISSN: 1064-8275, 1095-7197. DOI: `10.1137/18M1218686`. (Visited on 10/10/2023).

[38] Lukas Einkemmer, Alexander Ostermann, and Carmela Scalone. "A Robust and Conservative Dynamical Low-Rank Algorithm". In: *Journal of Computational Physics* 484 (July 2023), p. 112060. ISSN: 00219991. DOI: `10.1016/j.jcp.2023.112060`. (Visited on 10/23/2023).

[39] Francis Filbet and Tao Xiong. "Conservative Discontinuous Galerkin/Hermite Spectral Method for the Vlasov–Poisson System". In: *Communications on Applied Mathematics and Computation* 4.1 (Mar. 2022), pp. 34–59. ISSN: 2096-6385, 2661-8893. DOI: `10.1007/s42967-020-00089-z`. (Visited on 05/31/2022).

[40] Richard Fitzpatrick. *Plasma Physics: An Introduction.* Second edition. Boca Raton: CRC Press, 2022. ISBN: 978-1-03-220251-8 978-1-03-221407-8.

[41] M. Francisquez et al. "Conservative Discontinuous Galerkin Scheme of a Gyro-Averaged Dougherty Collision Operator". In: *Nuclear Fusion* 60.9 (Sept. 2020), p. 096021. ISSN: 0029-5515, 1741-4326. DOI: `10.1088/1741-4326/aba0c9`. (Visited on 11/18/2022).

[42] W Fundamenski. "Parallel Heat Flux Limits in the Tokamak Scrape-off Layer". In: *Plasma Physics and Controlled Fusion* 47.11 (Nov. 2005), R163–R208. ISSN: 0741-3335, 1361-6587. DOI: `10.1088/0741-3335/47/11/R01`. (Visited on 05/22/2023).

[43] Walter Gautschi. *Orthogonal Polynomials: Computation and Approximation.* Oxford University Press, Apr. 2004. ISBN: 978-0-19-850672-0 978-0-19-191657-1. DOI: `10.1093/oso/9780198506720.001.0001`. (Visited on 04/27/2023).

[44] R. J. Goldston and P. H. Rutherford. *Introduction to Plasma Physics.* Bristol, UK ; Philadelphia: Institute of Physics Pub, 1995. ISBN: 978-0-7503-0183-1.

[45] Sigal Gottlieb, David Ketcheson, and Chi-Wang Shu. *Strong Stability Preserving Runge-Kutta and Multistep Time Discretizations.* WORLD SCIENTIFIC, Jan. 2011. ISBN: 978-981-4289-26-9 978-981-4289-27-6. DOI: `10.1142/7498`. (Visited on 09/20/2023).

[46] Wei Guo, Jannatul Ferdous Ema, and Jing-Mei Qiu. "A Local Macroscopic Conservative (LoMaC) Low Rank Tensor Method with the Discontinuous Galerkin Method for the Vlasov Dynamics". In: *Communications on Applied Mathematics and Computation* (July 2023). ISSN: 2096-6385, 2661-8893. DOI: `10.1007/s42967-023-00277-7`. (Visited on 10/23/2023).

[47] Wei Guo and Jing-Mei Qiu. “A Conservative Low Rank Tensor Method for the Vlasov Dynamics”. In: (Jan. 2022). DOI: `10.48550/arXiv.2201.10397`. arXiv: `2201.10397`. (Visited on 05/09/2022).

[48] Wei Guo and Jing-Mei Qiu. “A Local Macroscopic Conservative (LoMaC) Low Rank Tensor Method for the Vlasov Dynamics”. In: (July 2022). DOI: `10.48550/arXiv.2207.00518`. arXiv: `2207.00518 [cs, math]`. (Visited on 10/23/2023).

[49] Jeffrey R. Haack, Cory D. Hauck, and Michael S. Murillo. “A Conservative, Entropic Multispecies BGK Model”. In: *Journal of Statistical Physics* 168.4 (Aug. 2017), pp. 826–856. ISSN: 0022-4715, 1572-9613. DOI: `10.1007/s10955-017-1824-9`. (Visited on 02/22/2022).

[50] Evan Habbershaw et al. *A Nonlinear, Conservative, Entropic Fokker-Planck Model for Multi-Species Collisions.* Apr. 2024. arXiv: `2404.11775 [math-ph]`. (Visited on 05/22/2024).

[51] Ammar Hakim and James Juno. “Alias-Free, Matrix-Free, and Quadrature-Free Discontinuous Galerkin Algorithms for (Plasma) Kinetic Equations”. In: *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis.* Atlanta, GA, USA: IEEE, Nov. 2020, pp. 1–15. ISBN: 978-1-72819-998-6. DOI: `10.1109/SC41405.2020.00077`. (Visited on 10/31/2022).

[52] Ammar Hakim, James Juno, and Gregory W. Hammett. “Conservative Discontinuous Galerkin Schemes for Nonlinear Dougherty–Fokker–Planck Collision Operators”. In: *Journal of Plasma Physics* 86.4 (Aug. 2020), p. 905860403. ISSN: 0022-3778, 1469-7807. DOI: `10.1017/S0022377820000586`. (Visited on 11/07/2021).

[53] Ammar H. Hakim. “Extended MHD Modelling with the Ten-Moment Equations”. In: *Journal of Fusion Energy* 27.1-2 (June 2008), pp. 36–43. ISSN: 0164-0313, 1572-9591. DOI: `10.1007/s10894-007-9116-z`. (Visited on 08/26/2022).

[54] G. W. Hammett, W. Dorland, and F. W. Perkins. “Fluid Models of Phase Mixing, Landau Damping, and Nonlinear Gyrokinetic Dynamics”. In: *Physics of Fluids B: Plasma Physics* 4.7 (Mar. 1992), pp. 2052–2061. ISSN: 0899-8221. DOI: `10.1063/1.860014`. (Visited on 08/31/2022).

[55] Gregory W. Hammett and Francis W. Perkins. “Fluid Moment Models for Landau Damping with Application to the Ion-Temperature-Gradient Instability”. In: *Physical Review Letters* 64.25 (June 1990), pp. 3019–3022. ISSN: 0031-9007. DOI: `10.1103/PhysRevLett.64.3019`. (Visited on 08/17/2022).

[56] R D Hazeltine. “Recursive Derivation of Drift-Kinetic Equation”. In: *Plasma Physics* 15.1 (Jan. 1973), pp. 77–80. ISSN: 0032-1028. DOI: `10.1088/0032-1028/15/1/009`. (Visited on 05/28/2024).

[57] R. D. Hazeltine, M. Kotschenreuther, and P. J. Morrison. “A Four-Field Model for Tokamak Plasma Dynamics”. In: *The Physics of Fluids* 28.8 (Aug. 1985), pp. 2466–2477. ISSN: 0031-9171. DOI: `10.1063/1.865255`. (Visited on 05/28/2024).

[58] Jan S. Hesthaven and Tim Warburton. *Nodal Discontinuous Galerkin Methods.* Ed. by J. E. Marsden, L. Sirovich, and S. S. Antman. Vol. 54. Texts in Applied Mathematics. New York, NY: Springer New York, 2008. ISBN: 978-0-387-72065-4 978-0-387-72067-8. DOI: `10.1007/978-0-387-72067-8`. (Visited on 01/17/2024).

[59] Andrew Ho, Iman Anwar Michael Datta, and Uri Shumlak. "Physics-Based-Adaptive Plasma Model for High-Fidelity Numerical Simulations". In: *Frontiers in Physics* 6 (Sept. 2018), p. 105. ISSN: 2296-424X. DOI: `10.3389/fphy.2018.00105`. (Visited on 01/31/2022).

[60] Thomas Y. Hou and Ruo Li. "Computing Nearly Singular Solutions Using Pseudo-Spectral Methods". In: *Journal of Computational Physics* 226.1 (Sept. 2007), pp. 379–397. ISSN: 00219991. DOI: `10.1016/j.jcp.2007.04.014`. (Visited on 09/18/2023).

[61] C. T. Hsu, R. D. Hazeltine, and P. J. Morrison. "A Generalized Reduced Fluid Model with Finite Ion-Gyroradius Effects". In: *The Physics of Fluids* 29.5 (May 1986), pp. 1480–1487. ISSN: 0031-9171. DOI: `10.1063/1.865665`. (Visited on 05/28/2024).

[62] Jingwei Hu and Yubo Wang. "An Adaptive Dynamical Low Rank Method for the Nonlinear Boltzmann Equation". In: *Journal of Scientific Computing* 92.2 (Aug. 2022), p. 75. ISSN: 0885-7474, 1573-7691. DOI: `10.1007/s10915-022-01934-4`. (Visited on 08/30/2022).

[63] J. D. Huba. "Finite Larmor Radius Magnetohydrodynamics of the Rayleigh–Taylor Instability". In: *Physics of Plasmas* 3.7 (July 1996), pp. 2523–2532. ISSN: 1070-664X, 1089-7674. DOI: `10.1063/1.871970`. (Visited on 01/21/2024).

[64] J. D. Huba. "The Kelvin-Helmholtz Instability: Finite Larmor Radius Magnetohydrodynamics". In: *Geophysical Research Letters* 23.21 (Oct. 1996), pp. 2907–2910. ISSN: 00948276. DOI: `10.1029/96GL02767`. (Visited on 08/29/2022).

[65] J. D. Huba, A. B. Hassam, and P. Satyanarayana. "Nonlocal Theory of the Rayleigh–Taylor Instability in the Limit of Unmagnetized Ions". In: *Physics of Fluids B: Plasma Physics* 1.4 (Apr. 1989), pp. 931–941. ISSN: 0899-8221. DOI: `10.1063/1.859017`. (Visited on 03/24/2024).

[66] J. D. Huba, J. G. Lyon, and A. B. Hassam. "Theory and Simulation of the Rayleigh-Taylor Instability in the Limit of Large Larmor Radius". In: *Physical Review Letters* 59.26 (Dec. 1987), pp. 2971–2974. ISSN: 0031-9007. DOI: `10.1103/PhysRevLett.59.2971`. (Visited on 01/21/2024).

[67] Shi Jin and Li Wang. "An Asymptotic Preserving Scheme for the Vlasov-Poisson-Fokker-Planck System in the High Field Regime". In: *Acta Mathematica Scientia* 31.6 (Nov. 2011), pp. 2219–2232. ISSN: 02529602. DOI: `10.1016/S0252-9602(11)60395-0`. (Visited on 07/02/2021).

[68] J. Juno et al. "Discontinuous Galerkin Algorithms for Fully Kinetic Plasmas". In: *Journal of Computational Physics* 353 (Jan. 2018), pp. 110–147. ISSN: 00219991. DOI: `10.1016/j.jcp.2017.10.009`. arXiv: `1705.05407`. (Visited on 02/02/2022).

[69] James Juno et al. *A New Hybrid: Kinetic Ions with Parallel-Kinetic-Perpendicular-Moment (Pkpm) Electrons*. Spokane, WA, Nov. 2023. (Visited on 05/28/2024).

[70] Emil Kieri, Christian Lubich, and Hanna Walach. "Discretized Dynamical Low-Rank Approximation in the Presence of Small Singular Values". In: *SIAM Journal on Numerical Analysis* 54.2 (Jan. 2016), pp. 1020–1038. ISSN: 0036-1429, 1095-7170. DOI: `10.1137/15M1026791`. (Visited on 06/06/2021).

[71] Othmar Koch and Christian Lubich. “Dynamical Low-Rank Approximation”. In: *SIAM Journal on Matrix Analysis and Applications* 29.2 (Jan. 2007), pp. 434–454. ISSN: 0895-4798, 1095-7162. DOI: `10.1137/050639703`. (Visited on 09/28/2021).

[72] Othmar Koch and Christian Lubich. “Dynamical Tensor Approximation”. In: *SIAM Journal on Matrix Analysis and Applications* 31.5 (Jan. 2010), pp. 2360–2375. ISSN: 0895-4798, 1095-7162. DOI: `10.1137/09076578X`. (Visited on 09/28/2021).

[73] Julian Koellermeier, Philipp Krah, and Jonas Kusch. *Macro-Micro Decomposition for Consistent and Conservative Model Order Reduction of Hyperbolic Shallow Water Moment Equations: A Study Using POD-Galerkin and Dynamical Low Rank Approximation.* Feb. 2023. arXiv: `2302.01391 [physics]`. (Visited on 02/16/2024).

[74] O. Koshkarov et al. “The Multi-Dimensional Hermite-discontinuous Galerkin Method for the Vlasov–Maxwell Equations”. In: *Computer Physics Communications* 264 (July 2021), p. 107866. ISSN: 00104655. DOI: `10.1016/j.cpc.2021.107866`. (Visited on 09/09/2022).

[75] Lev Landau. “THE TRANSPORT EQUATION IN THE CASE OF COULOMB INTERACTIONS”. In: *Collected Papers of L.D. Landau.* Elsevier, 1965, pp. 163–170. ISBN: 978-0-08-010586-4. DOI: `10.1016/B978-0-08-010586-4.50029-8`. (Visited on 01/13/2024).

[76] A. Lenard and Ira B. Bernstein. “Plasma Oscillations with Diffusion in Velocity Space”. In: *Physical Review* 112.5 (Dec. 1958), pp. 1456–1459. ISSN: 0031-899X. DOI: `10.1103/PhysRev.112.1456`. (Visited on 01/13/2024).

[77] Randall J. LeVeque. *Finite Volume Methods for Hyperbolic Problems.* 1st ed. Cambridge University Press, Aug. 2002. ISBN: 978-0-521-81087-6 978-0-521-00924-9 978-0-511-79125-3. DOI: `10.1017/CBO9780511791253`. (Visited on 05/08/2024).

[78] Randall J. LeVeque. *Numerical Methods for Conservation Laws.* 2nd ed. Lectures in Mathematics ETH Zürich. Basel ; Boston: Birkhäuser Verlag, 1992. ISBN: 978-3-7643-2723-1 978-0-8176-2723-2.

[79] Christian Lubich and Ivan V. Oseledets. “A Projector-Splitting Integrator for Dynamical Low-Rank Approximation”. In: *BIT Numerical Mathematics* 54.1 (Mar. 2014), pp. 171–188. ISSN: 0006-3835, 1572-9125. DOI: `10.1007/s10543-013-0454-0`. (Visited on 04/13/2022).

[80] Alan Macmahon. “Finite Gyro-Radius Corrections to the Hydromagnetic Equations for a Vlasov Plasma”. In: *The Physics of Fluids* 8.10 (Oct. 1965), pp. 1840–1845. ISSN: 0031-9171. DOI: `10.1063/1.1761116`. (Visited on 05/15/2024).

[81] E. T. Meier and U. Shumlak. “Development of Five-Moment Two-Fluid Modeling for Z-pinch Physics”. In: *Physics of Plasmas* 28.9 (Sept. 2021), p. 092512. ISSN: 1070-664X, 1089-7674. DOI: `10.1063/5.0058420`. (Visited on 10/09/2021).

[82] Barry Merriman. “Understanding the Shu–Osher Conservative Finite Difference Form”. In: *Journal of Scientific Computing* 19.1/3 (2003), pp. 309–322. ISSN: 08857474. DOI: `10.1023/A:1025312210724`. (Visited on 05/14/2024).

[83] S. T. Miller and U. Shumlak. “A Multi-Species 13-Moment Model for Moderately Collisional Plasmas”. In: *Physics of Plasmas* 23.8 (Aug. 2016), p. 082303. ISSN: 1070-664X, 1089-7674. DOI: `10.1063/1.4960041`. (Visited on 02/11/2022).

[84] Zhuogang Peng and Ryan G. McClarren. “A High-Order/Low-Order (HOLO) Algorithm for Preserving Conservation in Time-Dependent Low-Rank Transport Calculations”. In: *Journal of Computational Physics* 447 (Dec. 2021), p. 110672. ISSN: 00219991. DOI: `10.1016/j.jcp.2021.110672`. (Visited on 05/04/2023).

[85] F. Poupaud. “Runaway Phenomena and Fluid Approximation Under High Fields in Semiconductor Kinetic Theory”. In: *ZAMM - Journal of Applied Mathematics and Mechanics / Zeitschrift für Angewandte Mathematik und Mechanik* 72.8 (1992), pp. 359–372. ISSN: 00442267, 15214001. DOI: `10.1002/zamm.19920720813`. (Visited on 03/31/2022).

[86] J. J. Ramos. “Fluid Formalism for Collisionless Magnetized Plasmas”. In: *Physics of Plasmas* 12.5 (May 2005), p. 052102. ISSN: 1070-664X, 1089-7674. DOI: `10.1063/1.1884128`. (Visited on 05/15/2024).

[87] J. J. Ramos. “Fluid Theory of Magnetized Plasma Dynamics at Low Collisionality”. In: *Physics of Plasmas* 14.5 (May 2007), p. 052506. ISSN: 1070-664X, 1089-7674. DOI: `10.1063/1.2717595`. (Visited on 05/15/2024).

[88] A S Richardson. *2019 NRL PLASMA FORMULARY*. US Naval Research Laboratory, 2019.

[89] K. V. Roberts and J. B. Taylor. “Magnetohydrodynamic Equations for Finite Larmor Radius”. In: *Physical Review Letters* 8.5 (Mar. 1962), pp. 197–198. ISSN: 0031-9007. DOI: `10.1103/PhysRevLett.8.197`. (Visited on 05/20/2024).

[90] Youcef Saad and Martin H. Schultz. “GMRES: A Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear Systems”. In: *SIAM Journal on Scientific and Statistical Computing* 7.3 (July 1986), pp. 856–869. ISSN: 0196-5204, 2168-3417. DOI: `10.1137/0907058`. (Visited on 03/14/2022).

[91] Chi-Wang Shu. “Essentially Non-Oscillatory and Weighted Essentially Non-Oscillatory Schemes for Hyperbolic Conservation Laws”. In: *Advanced Numerical Approximation of Nonlinear Hyperbolic Equations*. Vol. 1697. Berlin, Heidelberg: Springer Berlin Heidelberg, 1998, pp. 325–432. ISBN: 978-3-540-64977-9 978-3-540-49804-9. DOI: `10.1007/BFb0096355`. (Visited on 09/19/2023).

[92] Chi-Wang Shu. “High Order Weighted Essentially Nonoscillatory Schemes for Convection Dominated Problems”. In: *SIAM Review* 51.1 (Feb. 2009), pp. 82–126. ISSN: 0036-1445, 1095-7200. DOI: `10.1137/070679065`. (Visited on 05/09/2023).

[93] Chi-Wang Shu and Stanley Osher. “Efficient Implementation of Essentially Non-Oscillatory Shock-Capturing Schemes”. In: *Journal of Computational Physics* 77.2 (Aug. 1988), pp. 439–471. ISSN: 00219991. DOI: `10.1016/0021-9991(88)90177-5`. (Visited on 05/20/2023).

[94] U. Shumlak et al. “Advanced Physics Calculations Using a Multi-Fluid Plasma Model”. In: *Computer Physics Communications* 182.9 (Sept. 2011), pp. 1767–1770. ISSN: 00104655. DOI: `10.1016/j.cpc.2010.12.048`. (Visited on 10/20/2020).

[95] Andrei N Simakov and Peter J Catto. "Momentum Transport in Arbitrary Mean-Free Path Plasma with a Maxwellian Lowest Order Distribution Function". In: *Plasma Physics and Controlled Fusion* 49.6 (June 2007), pp. 729–752. ISSN: 0741-3335, 1361-6587. DOI: `10.1088/0741-3335/49/6/004`. (Visited on 05/17/2024).

[96] B. Srinivasan and A. Hakim. "Role of Electron Inertia and Electron/Ion Finite Larmor Radius Effects in Low-Beta, Magneto-Rayleigh-Taylor Instability". In: *Physics of Plasmas* 25.9 (Sept. 2018), p. 092108. ISSN: 1070-664X, 1089-7674. DOI: `10.1063/1.5046098`. (Visited on 01/21/2024).

[97] Rich Sutton. *The Bitter Lesson*. Mar. 2019. (Visited on 06/12/2024).

[98] Yu Takagaki. "Development of a 5N-moment Multi-Fluid Plasma Model for DT Fusion in an Axisymmetric Z Pinch". PhD thesis. University of Washington, 2023.

[99] Lloyd N. Trefethen and David Bau. *Numerical Linear Algebra*. Philadelphia: Society for Industrial and Applied Mathematics, 1997. ISBN: 978-0-89871-361-9.

[100] Philipp Ulbl, Dominik Michels, and Frank Jenko. "Implementation and Verification of a Conservative, MULTI-SPECIES , GYRO-AVERAGED , Full- f , LENARD-BERNSTEIN /Dougherty Collision Operator in the Gyrokinetic Code GENE-X". In: *Contributions to Plasma Physics* (Dec. 2021). ISSN: 0863-1042, 1521-3986. DOI: `10.1002/ctpp.202100180`. (Visited on 02/22/2022).

[101] Takayuki Umeda, Satoshi Ueno, and Takuma K M Nakamura. "Ion Kinetic Effects on Nonlinear Processes of the Kelvin–Helmholtz Instability". In: *Plasma Physics and Controlled Fusion* 56.7 (July 2014), p. 075006. ISSN: 0741-3335, 1361-6587. DOI: `10.1088/0741-3335/56/7/075006`. (Visited on 09/26/2022).

[102] Takayuki Umeda et al. "Evaluating Gyro-Viscosity in the Kelvin-Helmholtz Instability by Kinetic Simulations". In: *Physics of Plasmas* 23.5 (May 2016), p. 054506. ISSN: 1070-664X, 1089-7674. DOI: `10.1063/1.4952632`. (Visited on 05/15/2024).

[103] Bram Van Leer. "Towards the Ultimate Conservative Difference Scheme. IV. A New Approach to Numerical Convection". In: *Journal of Computational Physics* 23.3 (Mar. 1977), pp. 276–299. ISSN: 00219991. DOI: `10.1016/0021-9991(77)90095-X`. (Visited on 05/08/2024).

[104] Bram Van Leer. "Towards the Ultimate Conservative Difference Scheme. V. A Second-Order Sequel to Godunov's Method". In: *Journal of Computational Physics* 32.1 (July 1979), pp. 101–136. ISSN: 00219991. DOI: `10.1016/0021-9991(79)90145-1`. (Visited on 05/08/2024).

[105] Juris Vencels et al. "SpectralPlasmaSolver: A Spectral Code for Multiscale Simulations of Collisionless, Magnetized Plasmas". In: *Journal of Physics: Conference Series* 719 (May 2016), p. 012022. ISSN: 1742-6588, 1742-6596. DOI: `10.1088/1742-6596/719/1/012022`. (Visited on 06/09/2022).

[106] Cédric Villani. "A Review of Mathematical Topics in Collisional Kinetic Theory". In: *Handbook of Mathematical Fluid Dynamics*. Vol. 1. Elsevier, 2002, pp. 71–74. ISBN: 978-0-444-50330-5. DOI: `10.1016/S1874-5792(02)80004-0`. (Visited on 04/13/2022).

[107] G. V. Vogman and J. H. Hammer. "High-Fidelity Kinetic Modeling of Instabilities and Gyromotion Physics in Nonuniform Low-Beta Plasmas". In: *Physics of Plasmas* 28.6 (June 2021), p. 062103. ISSN: 1070-664X, 1089-7674. DOI: `10.1063/5.0045983`. (Visited on 05/15/2024).

[108] G. V. Vogman et al. "Two-Fluid and Kinetic Transport Physics of Kelvin–Helmholtz Instabilities in Nonuniform Low-Beta Plasmas". In: *Physics of Plasmas* 27.10 (Oct. 2020), p. 102109. ISSN: 1070-664X, 1089-7674. DOI: `10.1063/5.0014489`. (Visited on 05/10/2021).

[109] G.V. Vogman, P. Colella, and U. Shumlak. "Dory–Guest–Harris Instability as a Benchmark for Continuum Kinetic Vlasov–Poisson Simulations of Magnetized Plasmas". In: *Journal of Computational Physics* 277 (Nov. 2014), pp. 101–120. ISSN: 00219991. DOI: `10.1016/j.jcp.2014.08.014`. (Visited on 05/02/2021).

[110] Linjin Zheng. "Perpendicular Magnetofluid Theory for Magnetically Confined Plasmas in the Collisionless Limit". In: *Physics of Plasmas* 27.5 (May 2020), p. 052503. ISSN: 1070-664X, 1089-7674. DOI: `10.1063/1.5144676`. (Visited on 05/15/2024).

Appendix

A.1 Linear analysis of electrostatic Vlasov oscillations

The electrostatic, collisionless, normalized Vlasov equation for a single species is

$$\partial_t f + \boldsymbol{v} \cdot \nabla_{\boldsymbol{x}} f + \omega_p \tau \frac{Z}{A} \boldsymbol{E} \cdot \nabla_{\boldsymbol{v}} f = 0. \tag{A.1.1}$$

To analyze this system, we suppose that it depends on a single space and a single velocity dimension, namely, that we are in a "1D1V" situation. We also adopt the convention $Z = A = 1$. Note that this is "flipped" from the likely physical interpretation of (A.1.1), which is that it describes the motion of an electron species against a static ion background. These assumptions further simplify our notation:

$$\partial_t f + v \partial_x f + \omega_p \tau E \partial_v f = 0. \tag{A.1.2}$$

This Vlasov equation is coupled with the normalized Gauss's law for the electric field, which states

$$\frac{\partial^2 \phi}{\partial x^2} = \omega_p \tau \left(n_0 - \int_{\mathbb{R}} f \, \mathrm{d}v \right), \quad E = -\partial_x \phi, \tag{A.1.3}$$

where n_0 is a fixed background density of ions.

To analyze (A.1.2) and (A.1.3), we begin by linearizing around an equilibrium distribution,

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

where f_1 is supposed small. Note that from spatial homogeneity of the equilibrium, $\phi = \phi_1$ and $E = E_1$. The linearization gives

$$\partial_t f_1 + v \partial_x f_1 + \omega_p \tau E_1 \partial_v f_0 = 0, \tag{A.1.4}$$

$$\frac{\partial^2 \phi_1}{\partial x^2} = \omega_p \tau \left(- \int_{\mathbb{R}} f_1 \, \mathrm{d}v \right). \tag{A.1.5}$$

We now suppose a Fourier normal mode ansatz for each perturbed quantity ψ_1:

$$\psi_1 = \tilde{\psi}_1 e^{ikx - i\omega t}. \tag{A.1.6}$$

Substituting (A.1.6) into (A.1.4) gives

$$-i\omega f_1 + ikvf_1 + \omega_p\tau(ik\phi)\partial_v f_0 = 0 \tag{A.1.7}$$

$$-k^2\phi_1 = -\omega_p\tau \int_{\mathbb{R}} f_1 \, \mathrm{d}v. \tag{A.1.8}$$

Solving (A.1.7) for f_1,

$$f_1 = \frac{-\omega_p\tau(ik\phi_1)\partial_v f_0}{i(kv - \omega)}.$$

We now substitute f_1 into (A.1.8) and move all terms to the left-hand side:

$$\phi_1 \left(1 + \frac{(\omega_p\tau)^2}{k^2} \int_{\mathbb{R}} \frac{\partial_v f_0}{v - \omega/k} \, \mathrm{d}v\right) = 0. \tag{A.1.9}$$

Equation (A.1.9) gives a compatibility condition for the existence of electrostatic oscillations. That is, for $\phi_1 = 0$, we must have

$$\epsilon(\omega, k) = 1 + \frac{(\omega_p\tau)^2}{k^2} \int_{\mathbb{R}} \frac{\partial_v f_0}{v - \omega/k} \, \mathrm{d}v = 0. \tag{A.1.10}$$

This is an implicit relation between ω and k which defines what is known as a dispersion relation. The function $\epsilon(\omega, k)$ is known as the dielectric function.

We should be careful to note that the formula given for $\epsilon(\omega, k)$ is defined only for $\mathrm{Im}(\omega) > 0$ due to the singularity at $v = \omega/k$. This case corresponds to exponential growth. For $\mathrm{Im}(\omega) \leq 0$, corresponding to undamped or damped waves, the dielectric function must be analytically continued into the lower half plane. Most introductory texts on plasma physics include a treatment of this procedure, which involves including a contribution from the residue of the integrand, which in turn modifies the zeros of the dielectric function, producing collisionless damping. See for example [40, 44].

For our purposes, we are interested in computing dispersion relations for the case where f_0 is a single Maxwellian or a superposition of Maxwellians, which is substantially aided by the introduction of the *plasma dispersion function* [22],

$$Z(\zeta) = \frac{1}{\sqrt{\pi}} \int_{\mathbb{R}} \frac{e^{-z^2}}{z - \zeta} \mathrm{d}z, \quad \mathrm{Im}(\zeta) > 0,$$

and defined by analytic continuation to the lower complex half-plane. Now, suppose that the equilibrium distribution function is given as the sum of Maxwellians $f_0 = M_1 + M_2 + \cdots$, each defined via their three macroscopic parameters n_i, u_i, v_{ti} as

$$M_i(v; n_i, u_i, v_{ti}) = \frac{n_i}{\sqrt{2\pi} v_{ti}} e^{-\frac{(v-u_i)^2}{2v_{ti}^2}}.$$

Defining $\xi_i = (v - u_i)/\sqrt{2} v_{ti}$,

$$\partial_v M_i(v) = -\sqrt{2}\xi_i M_i(v) = \frac{n_i}{\sqrt{\pi} v_{ti}} e^{-\xi_i^2}.$$

We therefore have

$$\int_{\mathbb{R}} \frac{\partial_v M_i(v)}{v - \omega/k} \, \mathrm{d}v = \frac{n_i}{\sqrt{\pi} v_{ti}} \int_{\mathbb{R}} \frac{\xi_i e^{-\xi_i^2}}{\xi_i + (u_i - \omega/k)/\sqrt{2} v_{ti}} \, \mathrm{d}\xi_i. \tag{A.1.11}$$

Defining the Doppler-shifted, normalized wave speed c_i by

$$c_i = \frac{\omega/k - u_i}{\sqrt{2} v_{ti}},$$

it is a matter of algebra to show

$$\int_{\mathbb{R}} \frac{\partial_v M_i(v)}{v - \omega/k} \, \mathrm{d}v = \frac{n_i}{\sqrt{\pi} v_{ti}} \left(1 + c_i Z(c_i)\right). \tag{A.1.12}$$

The dielectric function for the superposition of Maxwellians, then, is

$$\epsilon\left(\omega, k;\, f_0 = \sum_i M_i\right) = 1 + \frac{(\omega_p \tau)^2}{k^2} \sum_i \frac{n_i}{\sqrt{\pi} v_{ti}} [1 + c_i Z(c_i)].$$

To estimate zeros of $\epsilon(\omega, k)$, we take a page out of Daniel Crews's book [22] and plot contours of the real and imaginary parts in the complex plane, looking for intersections.