A Comparative Study of Rayleigh–Taylor Instability in Plasmas using Kinetic and Parallel Kinetic-Perpendicular Moment (PKPM) Models

Teja Naga Vardhan Peddi

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

Master of Science

University of Washington

2026

Committee:

Bhuvana Srinivasan, Chair

Uri Shumlak

Program Authorized to Offer Degree:
William E. Boeing Department of Aeronautics and Astronautics

University of Washington

Abstract

A Comparative Study of Rayleigh–Taylor Instability in Plasmas using Kinetic and Parallel Kinetic-Perpendicular Moment (PKPM) Models

Teja Naga Vardhan Peddi

Chair of the Supervisory Committee:
Professor Bhuvana Srinivasan
William E. Boeing Department of Aeronautics and Astronautics

Rayleigh-Taylor instability (RTI) is a fundamental instability that occurs when a denser medium is accelerated into a lighter one, and it plays an important role in many plasma systems, including laboratory plasmas, astrophysical flows, and fusion-related environments. While RTI has been studied extensively using fluid models and fully kinetic descriptions, it is less clear how well reduced kinetic-fluid models can reproduce the essential instability physics at a lower computational cost. This thesis examines that question by comparing fully kinetic and Parallel Kinetic-Perpendicular Moment (PKPM) simulations of RTI performed with the Gkeyll framework.

The work develops a consistent comparison between the two models using matched initial conditions, common boundary treatments, and contour-based diagnostics for interface evolution and growth-rate extraction. The study focuses on the evolution of density structure, instability growth, and transport behavior under different collisionalities and PKPM field-direction choices. For the Kn = 0.01 case, PKPM reproduces the main qualitative behavior of the kinetic reference, including the interface evolution and fitted growth rate, while showing greater sensitivity in the later nonlinear morphology. The results show that the retained kinetic direction in PKPM acts as an important control on the solution: in-plane field directions remain closest to the kinetic reference, while out-of-plane directions produce

broader and more fluid-like interface structures. At Kn = 0.1, the comparison becomes more transport-dominated, and the PKPM field direction more strongly affects whether the interface remains diffusive or develops a clearer RTI head.

Overall, the results show that PKPM provides a useful reduced description of RTI. It captures the dominant growth and morphology trends of the fully kinetic model while requiring significantly lower computational cost, making it a practical tool for studying larger and more expensive plasma-instability problems.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

I would first like to express my sincere gratitude to my advisor, Dr. Bhuvana Srinivasan, for her patience, encouragement, and continued faith in me throughout this journey. Her flexibility in allowing me to balance my work and research made this thesis possible. I cannot imagine how different my career path would have been without her support and the opportunity she gave me. She helped me grow, understand things more clearly, and keep moving forward even when the process felt difficult. I owe her a great deal, and I will always be deeply grateful for her guidance and support.

I would also like to thank Dr. Uri Shumlak for agreeing to serve on my committee and for supporting this work even during his sabbatical. I truly appreciate his time, encouragement, and willingness to help move this work forward.

I am also grateful to Jimmy Juno, John Rodman, and my labmates Lela, Lucio, and Vignesh for their guidance, technical help, and support through discussions, installations, questions, and computational resources, all of which made this work much smoother.

Finally, and most importantly, I would like to thank my family. They have always been there for me, reminding me that they would stand by me no matter what happens in life. Their support gave me the confidence to keep going. Becoming the first graduate in my family was not easy, especially given the financial challenges along the way, and I know this achievement belongs to them as much as it does to me. They have been the pillar of my confidence, and I will always be grateful for their love, sacrifice, and belief in me.

DEDICATION

To my family

Chapter 1

INTRODUCTION

Plasma physics is a field with applications spanning a wide range of areas, from studying fundamental physics to practical engineering systems. It seeks to understand how the ionized matter behaves when subjected to a wide range of electric and magnetic fields, how energy and momentum flow across many orders of magnitude in spatial and temporal scales and how collective effects give rise to complex and often unexpected dynamics. These same questions are important to naturally occurring plasmas like the solar wind and the Earth's magnetosphere and also to technologies ranging from fusion devices to space propulsion to plasma-based materials processing [10, 16].

1.1 Importance of plasmas in natural and laboratory systems

Plasma is commonly referred to as the fourth state of matter and consists of free electrons, ions, and, in many cases, neutral particles that collectively exhibit behavior over a wide range of spatial and temporal scales [10]. Unlike neutral gases, plasmas respond strongly to electric and magnetic fields and can support a broad variety of phenomena, including waves, instabilities, magnetic confinement, particle transport, heating, and radiation. Because of this rich dynamical behavior, plasma physics occupies an important place in both fundamental science and engineering.

Plasma physics is especially important because it connects a common set of physical ideas to systems that may appear very different on the surface. The same basic principles that describe naturally occurring plasmas, such as the solar wind and magnetospheric environments, also apply in laboratory settings where plasmas are created, controlled, and used for specific purposes. In that sense, plasma physics is both a fundamental and an applied science. It asks how charged particles and electromagnetic fields interact, and how those interactions affect transport, confinement, stability, and overall system behavior. Many

naturally occurring systems exist primarily in the plasma state. These include the solar wind, planetary magnetospheres, stellar atmospheres, and several diffuse astrophysical environments [16]. In such settings, plasma behavior governs the transport of momentum and energy, the generation and evolution of magnetic fields, and the emergence of large-scale structure. The same framework is therefore relevant across a remarkably broad range of conditions, from dense laboratory plasmas to weakly collisional and magnetized systems in space.

Plasmas are equally important in laboratory and technological applications. Controlled thermonuclear fusion is one of the most prominent examples, where plasma must be heated, confined, and sustained under extreme conditions in order to achieve energy-producing reactions [10]. Plasma is also central to many modern technologies, including semiconductor processing, ion propulsion, and nanoparticle production. As a result, plasma physics naturally draws from and contributes to many disciplines, including astronomy, solid-state physics, fluid dynamics, nonlinear dynamics, nuclear engineering, plasma chemistry, and plasma-material interactions.

In all of these settings, plasma behavior is shaped not only by external forcing but also by self-organized collective effects arising from particles, waves, fields, and material boundaries. One major consequence of this collective behavior is that plasmas are highly susceptible to disturbances that can grow, interact, and reorganize the system. Some disturbances decay, while others develop into instabilities that significantly alter transport, confinement, mixing, and overall system evolution. For this reason, understanding the origin and development of plasma instabilities is essential to both interpreting naturally occurring plasma phenomena and improving simulative capabilities in many laboratory plasma systems used to simulate stellar plasmas or achieve fusion. This motivates the need to study not only plasmas in equilibrium, but also the mechanisms through which that equilibrium can be disturbed.

1.2 Collective behavior and instabilities in plasmas

One of the defining features of a plasma is that its behavior is collective. Because electrons and ions move in response to electromagnetic fields, the motion of one part of a plasma can influence the motion of the system as a whole through long-range Coulomb forces, rather

than only through local collisions between nearest neighbors. As a result, plasmas can support a wide variety of behavior, including waves, oscillations, transport processes, and instabilities [10].

Not every disturbance in a plasma leads to a major change in the system. Some disturbances decay because of damping mechanisms such as collisions or kinetic effects like Landau damping, while others simply propagate as waves. However, some disturbances grow by drawing energy from the plasma itself, for example from density gradients, temperature gradients, imposed forcing, or an unstable equilibrium. When that occurs, the plasma becomes unstable. Such instability can strongly alter plasma structure and the transport of particles, momentum, and energy, making plasma evolution difficult to predict even when the governing equations are known. The details of that evolution depend strongly on the plasma regime. Factors such as collisionality, magnetization, and the relative importance of fluid and kinetic effects can all change how a disturbance grows and saturates. In weakly collisional and magnetized systems, for example, behavior along magnetic field lines can differ significantly from behavior across them, and even small departures from equilibrium can trigger microphysical processes that influence large-scale dynamics [16].

For this reason, the study of plasma instabilities is central to both fundamental plasma physics and many practical applications. In fusion-related systems, instabilities can degrade confinement, redistribute heat and particles in undesirable ways, and limit device performance. In space and astrophysical plasmas, instabilities can generate structure, mediate transport, and drive mixing over a wide range of scales. Understanding how and why these instabilities arise, and what physical mechanisms control their growth and saturation, is therefore essential to predicting plasma behavior.

What makes the subject especially challenging is that plasma instabilities rarely evolve in isolation. Their growth is affected by gradients, boundaries, imposed magnetic fields, collisionality, and other properties of the surrounding plasma. As a result, the same basic disturbance may behave very differently under different physical conditions [10, 16]. In many cases, instabilities ultimately determine whether a plasma configuration remains useful, breaks down, or transitions into a very different state.

Among the many instabilities that occur in plasmas, the Rayleigh–Taylor instability

is especially important because it lies at the intersection of classical fluid dynamics and distinctly plasma-specific effects. Its basic physical mechanism is simple, but its development can be altered substantially once transport, collisions, magnetic fields, and kinetic effects are taken into account. That makes it both physically important and a valuable benchmark for testing plasma models.

1.3 Rayleigh-Taylor instability and its relevance in plasma physics

Among the many instabilities that appear in fluid and plasma systems, the Rayleigh–Taylor instability (RTI) is one of the most fundamental. In its simplest form, RTI occurs when a denser medium is accelerated into a lighter one. Under these conditions, even a small disturbance at the interface can grow with time, causing the heavier material to penetrate downward in the form of spikes while the lighter material rises upward in the form of bubbles [17]. Although the basic physical picture is simple, the resulting evolution can become highly complex, especially once nonlinear effects begin to dominate.

RTI has been studied for many years in classical fluids, where it serves as a standard problem for understanding interfacial instability, mixing, and growth dynamics [17]. In plasma systems, however, the problem becomes richer. RTI is relevant in a range of environments, from astrophysical settings to laboratory plasmas, and its evolution can influence mixing, transport, and the redistribution of mass and energy [19, 20]. In high-energy-density and fusion-related contexts, these effects are especially important because instability growth can degrade confinement, distort interfaces, and alter the expected evolution of the system.

What makes plasma RTI particularly interesting is that it is not always governed by fluid behavior alone. In some regimes, fluid models provide a useful first description and recover the expected large-scale growth of the interface. In others, kinetic effects begin to matter. Finite collisionality, transport across the interface, non-Maxwellian features in the distribution function, and, more generally, departures from local equilibrium can modify both the growth rate and the structure of the instability [19]. Previous work has also shown that viscosity, resistivity, magnetic fields, and thermal conduction can all influence RTI evolution in plasma environments [21, 2, 25]. As a result, the behavior observed in a plasma RTI problem may differ substantially from the classical hydrodynamic picture, even when

the underlying instability mechanism remains the same.

RTI is also appealing from a modeling perspective because it combines a familiar macroscopic instability mechanism with physics that can become strongly model-dependent once the plasma leaves the ideal fluid limit. In a highly collisional regime, one expects behavior closer to classical hydrodynamic intuition. As collisionality decreases, however, particles can stream more freely across the interface, and transport begins to alter both the sharpness of the interface and the rate at which the disturbance develops. This makes RTI a particularly useful comparison problem: the instability mechanism remains recognizable, but its quantitative evolution can depend strongly on the modeling framework being used [19].

This is precisely where the present work is situated. The question is not only whether RTI develops, but how its development is represented when the same physical problem is viewed through different plasma descriptions. A fully kinetic model provides access to the underlying distribution function and the transport processes built into it, while a reduced model such as PKPM is designed to retain selected kinetic physics at lower computational cost. Comparing them in the context of RTI therefore provides a practical way to examine what is preserved, what is approximated, and where the two descriptions begin to separate.

For this reason, RTI provides a useful setting for comparing plasma models. It is simple enough to define clearly, but still rich enough to expose the roles of transport, collisionality, and model assumptions in a direct way. That balance makes it well suited to the goals of the present work, where the objective is not only to study the instability itself but also to examine how its evolution is represented in different computational descriptions of plasma dynamics.

1.4 Modeling challenges: kinetic and reduced descriptions

Once Rayleigh–Taylor instability is considered in a plasma setting, the question is no longer only whether the instability grows, but also how well a chosen model can represent that growth. In the simplest regimes, fluid descriptions can capture the large-scale behavior reasonably well and remain useful for building physical intuition. They are also far less expensive computationally than kinetic models. At the same time, plasma dynamics are

not always well described by fluid closures alone. When the plasma is weakly collisional, when transport across the interface becomes important, or when the particle distribution departs significantly from a local Maxwellian, a fluid picture may no longer be sufficient to capture the relevant physics [19].

A fully kinetic description retains this additional information by evolving the particle distribution function in phase space rather than only a small set of its moments. As a result, kinetic models can provide direct access to velocity-space structure, non-equilibrium transport, and departures from local Maxwellian behavior. In the context of RTI, this makes it possible to study how collisionality and transport modify the instability beyond the fluid limit. The drawback is that this level of fidelity comes at a substantial computational cost, especially as the number of spatial and velocity dimensions increases [16, 15].

This is what motivates reduced models. Their purpose is not simply to lower computational cost, but to do so while preserving the physics most relevant to the problem of interest. The Parallel-Kinetic-Perpendicular-Moment (PKPM) model is one such approach. It retains kinetic resolution along the magnetic-field direction while treating the perpendicular dynamics through reduced moments. If that reduction is done carefully, it offers the possibility of capturing important kinetic effects without paying the full cost of a fully kinetic calculation [16]. That tradeoff makes PKPM particularly relevant to the present study. It occupies a useful middle ground: richer than a conventional fluid model, but much more affordable than a fully kinetic description. At the same time, that middle ground raises an important question. A reduced model may reproduce some aspects of the instability well while misrepresenting others, and those differences may not be obvious in advance. Comparing PKPM with a fully kinetic model is therefore not merely a performance benchmark; it is also a way to identify which aspects of RTI are represented similarly in the two descriptions and which depend more strongly on the level of kinetic fidelity retained.

That comparison, however, is not entirely straightforward. Differences between two simulations do not necessarily arise only from the physics content of the models. They may also reflect how the initial equilibrium is constructed, how forcing is applied, how boundary conditions are imposed, or how transport is represented numerically. RTI is especially sensitive to such details because the instability develops from a background state that must

remain carefully balanced at the start of the simulation. A fair comparison therefore requires both models to begin from closely matched physical and numerical conditions. One of the central tasks of this thesis is to construct that comparison as consistently as possible, so that the later differences between the fully kinetic and PKPM results can be interpreted meaningfully.

1.5 Objectives and scope of the present work

So far we have established that Rayleigh–Taylor instability is an important problem in plasma systems, linking interface growth, transport, and mixing to the broader evolution of unstable plasma configurations. This discussion also shows that the way this evolution is modeled matters. Fully kinetic models can resolve the particle distribution function directly and therefore represent transport effects with high fidelity, but they do so at significant computational cost. Reduced models such as PKPM offer a more affordable alternative, but their ability to reproduce the essential instability physics under comparable conditions must be tested carefully. The primary motivation for the present work is therefore to examine this tradeoff between physical fidelity and computational tractability.

The main objective of this thesis is to compare the evolution of Rayleigh–Taylor instability in two plasma modeling frameworks implemented in Gkeyll: a fully kinetic description and the Parallel-Kinetic-Perpendicular-Moment (PKPM) model. The goal is not only to determine whether both models exhibit instability growth, but also to evaluate how closely their predictions agree when the simulations are constructed to represent the same physical problem as consistently as possible. In that sense, this thesis is both a physics study and a modeling study. It is concerned with the instability itself, but also with the practical question of how to compare two descriptions in a way that leads to meaningful conclusions.

A major part of the work is therefore devoted to building a consistent comparison framework. In principle, that sounds straightforward: initialize the same Rayleigh–Taylor configuration in both models, evolve each system, and compare the results. In practice, however, the problem is more delicate. A meaningful comparison requires careful attention to the initial equilibrium, the imposed acceleration, the perturbation used to seed the instability, the role of collisions, and the boundary conditions. If the background state is not properly

balanced, or if the two models interpret the setup differently, then the resulting differences may reflect inconsistencies in setup rather than actual differences in model physics. One goal of this thesis is to make that distinction as clear as possible.

More specifically, the thesis addresses four main questions. First, does PKPM reproduce the main qualitative features of the instability seen in the fully kinetic simulations, including the evolution of the interface and the overall morphology of the density field? Second, do the two models predict comparable growth rates under the same physical setup? Third, how does the predicted behavior depend on the direction chosen for the retained kinetic response in PKPM? Fourth, how do the ideal and non-ideal transport terms evolve, and what do they reveal about how the reduced model represents the underlying transport physics?

A second theme of the thesis concerns the role of collisional modeling and parameter variation. In the kinetic framework, we examine how different collision operators, namely BGK and LBO, affect RTI growth and interface morphology. In the PKPM framework, we examine how the instability changes with Knudsen number and field-direction choice. This helps separate differences that arise from the reduced nature of PKPM from those that arise from the particular collisional model or parameter regime being used.

The scope of the present study is intentionally limited to a controlled comparison problem rather than a broad survey of all possible plasma RTI configurations. The goal is not to develop a complete theory of plasma RTI under every relevant physical effect, but to use a carefully constructed side-by-side comparison between a fully kinetic description and the reduced PKPM framework to gain insight into RTI under well-defined conditions. Several topics therefore remain outside the scope of this work. We do not attempt to cover the full range of magnetic-field geometries, collisional closures, or plasma parameter regimes relevant to RTI. While broader plasma applications motivate the present study, the simulations considered here are idealized and are intended primarily as a comparative modeling framework rather than as a direct representation of a specific experimental or astrophysical system.

Overall, the goal of this thesis is to determine what can be learned by studying Rayleigh–Taylor instability with both a high-fidelity kinetic approach and a reduced PKPM description under matched conditions. By comparing instability morphology, growth rates, transport effects,

collision-operator dependence, parameter sensitivity, and computational cost, we aim to identify the regimes in which the reduced model remains physically useful, the regimes in which important deviations begin to appear, and the considerations that should guide the interpretation of such comparisons.

1.6 Organization of the thesis

The remainder of this thesis is organized as follows. Chapter 2 presents the modeling framework and numerical methodology used in this work. It introduces the fully kinetic and PKPM descriptions, outlines the computational framework used for the simulations, and discusses the initialization strategy, boundary conditions, diagnostics, and other practical considerations required to construct comparable runs. Chapter 3 presents the simulation results and compares the outcomes of the different modeling approaches, with emphasis on instability evolution, growth behavior, transport effects, and the influence of collisionality and PKPM field direction. Finally, Chapter 4 summarizes the main findings of the thesis and outlines directions for future work.

Chapter 2

MODELING FRAMEWORK AND NUMERICAL METHODOLOGY

2.1 Hierarchy of plasma descriptions

The level of description required to capture plasma phenomena depends strongly on the physics of interest and on the computational cost that can be afforded. At one end of the hierarchy are fluid and moment models, which evolve macroscopic quantities such as density, bulk velocity, and pressure. At the other end are fully kinetic models, which evolve the particle distribution function in phase space. Between these two limits lie reduced descriptions that retain selected kinetic effects while avoiding the full cost of a high-dimensional phase-space calculation [10, 16, 15].

These descriptions are connected through velocity moments of the distribution function. For a species s with distribution function $f_s(\mathbf{x}, \mathbf{v}, t)$, the number density is

$$n_s = \int f_s \, d^3v, \tag{2.1}$$

the bulk velocity is

$$\mathbf{u}_s = \frac{1}{n_s} \int \mathbf{v} f_s \, d^3v, \tag{2.2}$$

Fluid models evolve some subset of these moments and require closure assumptions for the higher-order moments that are not retained explicitly. Fully kinetic models avoid this closure problem by evolving the distribution function itself, while reduced models aim to preserve the most important kinetic effects in a more computationally affordable form.

In this work, we compare two models that lie at different levels of this hierarchy: a fully kinetic description and the Parallel-Kinetic-Perpendicular-Moment (PKPM) model. The fully kinetic model serves as the higher-fidelity reference, while PKPM provides a reduced description designed to retain part of the kinetic physics at lower cost [16]. The following sections describe these two frameworks and the numerical methodology used to compare them in the context of Rayleigh–Taylor instability.

2.2 Fully kinetic description

2.2.1 Phase-space description of the plasma

In a fully kinetic description, the plasma is represented by its distribution function rather than by a small set of macroscopic variables alone. For every species s, the distribution function $f_s(\mathbf{x}, \mathbf{v}, t)$ gives the phase-space density of particles at position $\mathbf{x}$ with velocity $\mathbf{v}$ at time t. Macroscopic quantities such as density, bulk velocity, and pressure are then obtained as moments of the evolving distribution function [10, 15].

This description is especially valuable when the plasma is weakly collisional, when velocity-space structure matters, or when the distribution function departs significantly from a local Maxwellian. In such cases, low-order fluid closures may not capture the relevant transport and non-equilibrium physics with sufficient accuracy. A fully kinetic model retains this information directly, at the cost of resolving phase space rather than configuration space alone [15, 19].

More generally, a kinetic model may involve up to three configuration-space coordinates, three velocity-space coordinates, and time. The exact number of resolved dimensions depends on the problem being studied, but the defining feature of the kinetic approach is that it evolves the distribution function directly rather than only moments derived from it.

2.2.2 Governing kinetic equation

The evolution of the distribution function is governed by the Boltzmann equation. For a species s, it may be written as

$$\frac{\partial f_s}{\partial t} + \mathbf{v} \cdot \nabla_{\mathbf{x}} f_s + \frac{q_s}{m_s} (\mathbf{E} + \mathbf{v} \times \mathbf{B}) \cdot \nabla_{\mathbf{v}} f_s + \mathbf{a}_{\text{ext}} \cdot \nabla_{\mathbf{v}} f_s = C_s[f], \tag{2.3}$$

where q_s and m_s are the charge and mass of species s, $\mathbf{E}$ and $\mathbf{B}$ are the electric and magnetic fields, $\mathbf{a}_{\text{ext}}$ is any prescribed external acceleration, and $C_s[f]$ is a collision operator [10, 15].

This equation ((2.3)) expresses conservation of particles in phase space. The first term represents local time evolution, the second describes advection in configuration space, and the third and fourth terms describe acceleration in velocity space due to electromagnetic

forcing and any externally applied body force. The right-hand side accounts for collisional relaxation and velocity-space diffusion.

In the collisionless limit, $C_s[f] = 0$, the equation reduces to the Vlasov equation. For the body-force driven Rayleigh–Taylor problems considered in this thesis, we use a neutral-species form of the kinetic equation,

$$\frac{\partial f_s}{\partial t} + \mathbf{v} \cdot \nabla_{\mathbf{x}} f_s + \mathbf{a} \cdot \nabla_{\mathbf{v}} f_s = C_s[f], \tag{2.4}$$

where $\mathbf{a}$ is the imposed acceleration that drives the instability. In this setting, the Lorentz-force term does not appear explicitly because the kinetic reference is used in a neutral, body-force driven form rather than as a fully electromagnetic Vlasov–Maxwell system.

When self-consistent electromagnetic evolution is retained, the kinetic equation is coupled to Maxwell's equations through

$$\nabla \cdot \mathbf{E} = \frac{\rho_c}{\varepsilon_0}, \tag{2.5}$$

$$\nabla \cdot \mathbf{B} = 0, \tag{2.6}$$

$$\frac{\partial \mathbf{B}}{\partial t} = -\nabla \times \mathbf{E}, \tag{2.7}$$

$$\varepsilon_0 \frac{\partial \mathbf{E}}{\partial t} = \frac{1}{\mu_0} \nabla \times \mathbf{B} - \mathbf{J}, \tag{2.8}$$

where the charge density ρ_c and current density $\mathbf{J}$ from moments of the distribution function,

$$\rho_c = \sum_s q_s \int f_s \, d^3v, \tag{2.9}$$

$$\mathbf{J} = \sum_s q_s \int \mathbf{v} f_s \, d^3v. \tag{2.10}$$

In the present thesis, however, the fully kinetic RTI problem is used primarily in the body-force driven setting.

2.2.3 Velocity moments and macroscopic quantities

While the full kinetic simulation updates the distribution function directly, the standard macroscopic quantities are still recovered through velocity moments. The zeroth moment gives the number density,

$$n_s(\mathbf{x}, t) = \int f_s(\mathbf{x}, \mathbf{v}, t) \, d^3v, \tag{2.11}$$

and the first moment gives the bulk velocity,

$$\mathbf{u}_s(\mathbf{x},t) = \frac{1}{n_s}\int \mathbf{v} f_s(\mathbf{x},\mathbf{v},t)\, d^3v. \tag{2.12}$$

The mass density is then

$$\rho_s = m_s n_s. \tag{2.13}$$

To define higher moments, it is convenient to introduce the peculiar velocity

$$\mathbf{w} = \mathbf{v} - \mathbf{u}_s. \tag{2.14}$$

The second central moment gives the pressure tensor,

$$P_{s,ij} = m_s \int w_i w_j f_s\, d^3v, \tag{2.15}$$

which measures the spread of particle velocities about the bulk flow. The scalar pressure is obtained from the trace,

$$p_s = \frac{1}{3} P_{s,ii}, \tag{2.16}$$

and in normalized units the temperature is commonly written as

$$T_s = \frac{p_s}{n_s}. \tag{2.17}$$

It is important to note the distinction between the raw second moment tensor and the central second moment tensor. The raw second moment tensor is given by

$$M_{s,ij} = m_s \int v_i v_j f_s\, d^3v, \tag{2.18}$$

which may be decomposed as

$$M_{s,ij} = P_{s,ij} + \rho_s u_{s,i} u_{s,j}. \tag{2.19}$$

We see that this expression involves both the thermal motion via $P_{s,ij}$, and ordered bulk motion of the particles, via u_s.

The third central moment gives the heat-flux tensor,

$$Q_{s,ijk} = m_s \int w_i w_j w_k f_s\, d^3v, \tag{2.20}$$

and the usual heat-flux vector is obtained by contraction,

$$q_{s,k} = \frac{1}{2} Q_{s,iik} = \frac{m_s}{2} \int w^2 w_k f_s \, d^3v. \tag{2.21}$$

These higher moments are important because they quantify transport effects that are not captured solely through density, flow velocity, and scalar pressure.

The non-ideal stress tensor is defined for subsequent analysis as

$$\Pi_{s,ij} = P_{s,ij} - p_s \delta_{ij}, \tag{2.22}$$

emphasizing the anisotropic part of the stress. In this framework, the energy flux may be written as

$$\mathcal{Q}_{s,k} = \left(\frac{1}{2} \rho_s u_s^2 + \frac{5}{2} p_s \right) u_{s,k} + u_{s,i} \Pi_{s,ik} + q_{s,k}, \tag{2.23}$$

where

$$\mathcal{Q}_{s,k} = \int \frac{1}{2} m_s v^2 v_k f_s \, d^3v. \tag{2.24}$$

This form is useful because it separates ideal convective transport from heat-flux and non-ideal stress contributions [19].

2.2.4 Collision operators

The collision operator in Eq. (2.3) determines how the distribution relaxes and how collisional transport is represented. The collision term is modeled using simplified velocity-space operators rather than the full Landau operator. In this work, two collision models are relevant: the Bhatnagar–Gross–Krook (BGK) operator and the Lenard–Bernstein–O'Neil / Dougherty-type Fokker–Planck operator, referred to here as LBO.

Bhatnagar-Gross-Krook (BGK) operator is given as,

$$C_{\mathrm{BGK}}[f_s] = \nu_s \left(f_{M,s} - f_s \right), \tag{2.25}$$

where ν_s is the collision frequency and $f_{M,s}$ is a Maxwellian constructed from the local moments of the evolving distribution,

$$f_{M,s} = n_s \left(\frac{m_s}{2\pi T_s} \right)^{3/2} \exp \left[- \frac{m_s (\mathbf{v} - \mathbf{u}_s)^2}{2 T_s} \right]. \tag{2.26}$$

The BGK operator replaces the full collision term with a relaxation model that drives the distribution toward the local Maxwellian while preserving density, momentum, and energy when implemented consistently. It is simple, computationally efficient, and useful for controlled comparisons, especially in neutral-species problems [3, 19].

The LBO operator provides a different reduced collisional model by combining a drag term with a velocity-space diffusion term. In compact form, it may be written as

$$C_s[f] = \sum_r \nu_{sr} \nabla_{\mathbf{v}} \cdot \left[(\mathbf{v} - \mathbf{u}_{sr}) f_s + v_{\mathrm{th},sr}^2 \nabla_{\mathbf{v}} f_s \right], \tag{2.27}$$

where ν_{sr} is a collision frequency, $\mathbf{u}_{sr}$ is an effective drift velocity, and $v_{\mathrm{th},sr}$ is an effective thermal speed [18, 6, 12]. This form preserves the conservative structure needed for continuum DG implementations and is the main collisional operator used for the matched kinetic–PKPM comparisons in this thesis [12].

In the RTI studies most closely related to the present work, both BGK and LBO are used to represent finite collisionality. Their use makes it possible to examine how instability growth changes as the system moves between more kinetic and more fluid-like limits [19].

2.2.5 Continuum kinetic formulation in Gkeyll

The fully kinetic simulations in this thesis are carried out using the Gkeyll framework. Gkeyll solves continuum kinetic equations by discretizing the distribution function on a phase-space mesh rather than representing the plasma with macro-particles. The phase-space discretization is based on the discontinuous Galerkin (DG) method, which provides high-order accuracy together with a formulation well suited for advection-dominated kinetic equations. [15, 8].

In the present work, the fully kinetic model serves as the higher-fidelity reference for comparison with PKPM. Its main strength is that it evolves the full phase-space distribution function and therefore retains kinetic transport and higher moments without requiring a low-order closure from the outset. The details of the spatial discretization, basis choice, and simulation setup are presented in the following sections; for now, the important point is that the fully kinetic model provides the reference description against which the reduced PKPM model is evaluated.

2.3 Parallel-Kinetic-Perpendicular-Moment model

2.3.1 Motivation and physical setting

The Parallel-Kinetic-Perpendicular-Moment (PKPM) model is intended for weakly collisional, magnetized plasmas in which the dynamics parallel and perpendicular to the magnetic field can differ substantially. In such regimes, plasma particles moving along the magnetic field are mostly unhindered, whereas movement perpendicular to the field lines is restricted. So momentum and heat can evolve anisotropically. This separation of parallel and perpendicular dynamics motivates the use of reduced models that retain the most important kinetic physics along the field while reducing the perpendicular degrees of freedom.

The basic idea of PKPM is to exploit that anisotropy directly. Instead of evolving the full distribution function in all velocity-space directions, the model retains a kinetic description parallel to the magnetic field and represents the perpendicular dynamics through a reduced spectral or moment-based structure. In that sense, PKPM lies between fluid models and fully kinetic models: it retains more kinetic content than a low-order fluid closure, but at a cost that is much lower than that of solving the full Vlasov equation in all phase-space dimensions.

An important point given in Juno et al., is that the PKPM model used here is not introduced as an asymptotic expansion in a small parameter. Instead, it is constructed through coordinate transformations and a reduced representation in the perpendicular velocity. This gives the model a natural connection to magnetized plasma dynamics while keeping the reduction flexible enough for problems beyond the strict limits of traditional asymptotic models.

2.3.2 Field-aligned velocity coordinates

The construction of PKPM begins by decomposing the velocity $\mathbf{v}$ of a particle into components parallel ($v_\parallel$) and perpendicular ($v_\perp$) to the magnetic field. Let $\mathbf{b} = \mathbf{B}/|\mathbf{B}|$ be the magnetic-field unit vector, let $\boldsymbol{\tau}_1$ and $\boldsymbol{\tau}_2$ form a field-aligned basis spanning the plane perpendicular to $\mathbf{b}$. The velocity may then be written in field-aligned form $(\mathbf{b}, \boldsymbol{\tau}_1, \boldsymbol{\tau}_2)$ as

$$\mathbf{v} = v_{\parallel}\mathbf{b} + v_{\perp}\cos\theta\,\boldsymbol{\tau}_1 + v_{\perp}\sin\theta\,\boldsymbol{\tau}_2, \tag{2.28}$$

where $v_{\parallel}$ is the parallel velocity, $v_{\perp}$ is the perpendicular speed, and θ is the gyrophase angle.

This decomposition makes it possible to preserve full kinetic resolution only in the parallel direction while representing the perpendicular dependence with a reduced set of coefficients. The resulting model therefore avoids a uniform discretization of the full velocity space while still retaining the part of the kinetic response most directly aligned with the magnetic field.

2.3.3 Spectral reduction / model structure

After transforming to field-aligned coordinates, the perpendicular dependence of the distribution function is represented through a reduced spectral expansion. In practice, this means that PKPM replaces the full perpendicular velocity-space resolution with a smaller set of coefficients that retain the dominant perpendicular physics. The reduction is most effective when the perpendicular structure remains close enough to a low-order representation, while the parallel direction is left fully kinetic so that nontrivial field-aligned phase-space structure can still be resolved.

For this, the gyrotropic portion of the distribution function is expanded in Laguerre polynomials of the perpendicular speed $v_{\perp}$, as follows:

$$\bar{f}_0(\mathbf{x}, v_{\parallel}, v_{\perp}, t) = \sum_{k=0}^{\infty} F_k(\mathbf{x}, v_{\parallel}, t)\, G_M(v_{\perp}, T_{\perp})\, L_k\left(\frac{m v_{\perp}^2}{2T_{\perp}}\right), \tag{2.29}$$

where L_k denotes the Laguerre polynomials and

$$G_M(v_{\perp}, T_{\perp}) = \frac{m}{2\pi T_{\perp}} \exp\left(-\frac{m v_{\perp}^2}{2T_{\perp}}\right) \tag{2.30}$$

is the perpendicular Maxwellian weight.

In this approach, it is clear what the reduction means. The parallel speed is an explicit coordinate, allowing us to display streaming and the parallel part of phase-space structure. All of the perpendicular dependence is instead contained within the coefficients F_k. When

few such coefficients suffice, this is a six-dimensional kinetic problem that reduces to a smaller system of four-dimensional equations. The result is a hybrid description: kinetic in the parallel direction and reduced in the perpendicular one. This is the key modeling idea behind PKPM and is the reason the model can retain important anisotropic transport physics at significantly lower cost than a fully kinetic treatment.

2.3.4 Lowest-order PKPM variables

In order for PKPM to be used for practical computation, it is most frequently used in its lowest order. This is the simplest level of the original PKPM formulation and is obtained by keeping only the zeroth harmonic of the gyrophase Fourier expansion. In addition, we keep the first two Laguerre coefficients in perpendicular velocity, all higher-order harmonics/coefficients are set to zero. The kinetic quantities kept in this lowest-order are $F_0(\mathbf{x}, v_\parallel, t)$ and $G(\mathbf{x}, v_\parallel, t)$, where

$$G \equiv \frac{T_\perp}{m}(F_0 - F_1). \tag{2.31}$$

This definition is useful because it allows the perpendicular pressure to be evolved directly without introducing a separate evolution equation for $T_\perp$. In other words, the temperature normalization appearing in the Laguerre expansion (from Eq. 2.29) is absorbed into the evolution of G, which makes the reduced system more compact and easier to use in practice.

The main macroscopic quantities needed for closure are derived from moments of F_0 and G. The number density is

$$n = \int_{-\infty}^{\infty} F_0 \, dv_\parallel, \tag{2.32}$$

the parallel pressure is

$$p_\parallel = \int_{-\infty}^{\infty} m v_\parallel^2 F_0 \, dv_\parallel, \tag{2.33}$$

and the perpendicular pressure is

$$p_\perp = m \int_{-\infty}^{\infty} G \, dv_\parallel. \tag{2.34}$$

Similarly, the parallel and perpendicular heat-flux contributions are

$$q_\parallel = \int_{-\infty}^{\infty} \frac{1}{2} m v_\parallel^3 F_0 \, dv_\parallel, \tag{2.35}$$

and

$$q_\perp = T_\perp \int_{-\infty}^{\infty} v_\parallel (F_0 - F_1)\, dv_\parallel. \tag{2.36}$$

These relations show that we do have access to the anisotropy in the pressure and heat-flux through the kinetic variables.

2.3.5 Lowest-order PKPM equations

With just the Fourier zero mode and the first two Laguerre coefficients taken, the lowest-order PKPM model is made of two coupled 4-D kinetic equations for F_0 and (G).

The first governs F_0,

$$\begin{aligned} &\frac{\partial F_0}{\partial t} + \nabla_{\mathbf{x}} \cdot \left[(v_\parallel \mathbf{b} + \mathbf{u}) F_0 \right] \\ &\quad + \frac{\partial}{\partial v_\parallel} \left[\mathbf{b} \cdot \left(\frac{1}{\rho} \nabla_{\mathbf{x}} \cdot \mathbf{P} - v_\parallel \mathbf{b} \cdot \nabla_{\mathbf{x}} \mathbf{u} \right) F_0 + G \nabla_{\mathbf{x}} \cdot \mathbf{b} \right] = 0, \end{aligned} \tag{2.37}$$

while the second governs G,

$$\begin{aligned} &\frac{\partial G}{\partial t} + \nabla_{\mathbf{x}} \cdot \left[(v_\parallel \mathbf{b} + \mathbf{u}) G \right] \\ &\quad + \frac{\partial}{\partial v_\parallel} \left[\mathbf{b} \cdot \left(\frac{1}{\rho} \nabla_{\mathbf{x}} \cdot \mathbf{P} - v_\parallel \mathbf{b} \cdot \nabla_{\mathbf{x}} \mathbf{u} \right) G \right. \\ &\qquad\qquad \left. + \left(4G - 2 \frac{T_\perp}{m} F_0 \right) \frac{T_\perp}{m} \nabla_{\mathbf{x}} \cdot \mathbf{b} \right] \\ &\quad = \left[\mathbf{b} \cdot (\mathbf{b} \cdot \nabla_{\mathbf{x}} \mathbf{u}) - \nabla_{\mathbf{x}} \cdot (v_\parallel \mathbf{b} + \mathbf{u}) \right] G. \end{aligned} \tag{2.38}$$

Here, the advection terms $\nabla_{\mathbf{x}} \cdot [(v_\parallel \mathbf{b} + \mathbf{u}) F_0]$ and $\nabla_{\mathbf{x}} \cdot [(v_\parallel \mathbf{b} + \mathbf{u}) G]$ describe advection of F_0 and G, both in the sense of streaming along $\mathbf{b}$ and of bulk flow driven by $\mathbf{u}$. The velocity-space derivative terms describe acceleration along the magnetic field due to the forces arising from pressure gradients, spatial gradients of flow velocity and magnetic field curvature or torsion. Specifically, $\nabla_{\mathbf{x}} \cdot \mathbf{b}$ is coupling the magnetic geometry with the reduced perpendicular dynamics and is central to the unique features of the PKPM model. The important point is that the problem remains kinetic along the magnetic field line, as both distribution functions depend explicitly on $v_\parallel$ and yet the $v_\perp$ dependence is not fully resolved in the distribution. The structure in that direction enters only implicitly through the chosen Laguerre modes and the calculated moments thereof.

2.3.6 Momentum equation and closure

The PKPM system closes through the momentum equation for the bulk flow:

$$\frac{\partial(\rho \mathbf{u})}{\partial t} + \nabla_{\mathbf{x}} \cdot (\rho\, \mathbf{u} \otimes \mathbf{u} + \mathbf{P}) = \rho \frac{q}{m} \left(\mathbf{E} + \mathbf{u} \times \mathbf{B} \right), \tag{2.39}$$

with a pressure tensor of the form:

$$\mathbf{P} = p_{\parallel}\, \mathbf{b} \otimes \mathbf{b} + p_{\perp} \left(\mathbf{I} - \mathbf{b} \otimes \mathbf{b} \right). \tag{2.40}$$

Note that this point should be reiterated for PKPM because it does not simply prescribe the anisotropic pressure as a closure. Instead, $p_{\parallel}$ and $p_{\perp}$ self-consistently arise from the moments of the kinetic variables retained via Eqs. (2.33) and (2.34), meaning PKPM remains more tightly connected to the distribution function than a conventional low-order fluid closure.

When considering the full electromagnetic regime, the equations of PKPM are coupled to Maxwell's equations:

$$\frac{\partial \mathbf{B}}{\partial t} + \nabla_{\mathbf{x}} \times \mathbf{E} = 0, \tag{2.41}$$

$$\varepsilon_0 \mu_0 \frac{\partial \mathbf{E}}{\partial t} - \nabla_{\mathbf{x}} \times \mathbf{B} = -\mu_0 \mathbf{J}, \tag{2.42}$$

$$\nabla_{\mathbf{x}} \cdot \mathbf{E} = \frac{\rho_c}{\varepsilon_0}, \tag{2.43}$$

$$\nabla_{\mathbf{x}} \cdot \mathbf{B} = 0, \tag{2.44}$$

with the current density $\mathbf{J}$ given by

$$\mathbf{J} = \sum_s \frac{q_s}{m_s} \rho_s \mathbf{u}_s. \tag{2.45}$$

This demonstrates that PKPM is not simply a prescribed-field description, it constitutes a self-consistent kinetic-fluid electromagnetic model.

2.3.7 Relevance to the present work

PKPM is the reduced model that is being compared against the fully kinetic model in this thesis. Its value lies in the balance it offers between model fidelity and computational cost.

The model remains kinetic in the parallel direction, supports anisotropic pressure dynamics, and retains a close connection to distribution-function moments, while avoiding the cost of resolving the full perpendicular velocity space explicitly. That balance is precisely what makes PKPM central to the present study. The aim is not only to determine whether PKPM is computationally cheaper, but to assess whether it preserves the aspects of Rayleigh–Taylor instability that matter physically under conditions comparable to the fully kinetic model. If it does, then it becomes a promising tool for larger and more expensive studies. If important differences appear, then the comparison helps identify where in the process the reduced representation begins to lose physical information pertinent to the dynamics of the instability.

2.4 Computational framework

2.4.1 Gkeyll as the solver framework

Both the fully kinetic and PKPM simulations are carried out within the `Gkeyll` simulation framework developed at Princeton Plasma Physics Laboratory as a continuum plasma code capable of solving fully kinetic equations, multi-moment models, and reduced kinetic descriptions under a unified numerical framework. This is important because it allows the comparison to be made within a shared numerical environment rather than across unrelated codes. Both approaches use discontinuous Galerkin spatial discretization, explicit Runge–Kutta time integration, the same configuration-space geometry, and similar diagnostic output structures [8, 15, 19]. As a result, differences observed later in the thesis can be interpreted primarily as model differences rather than artifacts of unrelated numerical implementations.

In Gkeyll, the fully kinetic and PKPM models are advanced through different application layers `Vlasov.App` and `PKPM.App`, but both remain part of the same overall computational framework. This provides a common basis for domain decomposition, polynomial representation within each cell, explicit time stepping, and post-processing.

2.4.2 Domain geometry and orientation of the instability

The computational domain is two-dimensional in configuration space and is taken as

$$(x, y) \in [-L_x, L_x] \times [-L_y, L_y], \tag{2.46}$$

with

$$L_x = 1.5, \qquad L_y = 0.75 \tag{2.47}$$

in the reference setup. Both models use the same configuration-space resolution,

$$N_x = N_y = 64, \tag{2.48}$$

so that the comparison is carried out on a common spatial grid.

In the present implementation, the Rayleigh–Taylor instability develops along the x-direction rather than the vertical direction used in many standard presentations of the problem. This rotation is made for computational reasons, not physical ones. In the Gkeyll workflow used here, the fixed-function boundary treatment needed to maintain reservoir-like behavior is applied in the x-direction, while the y-direction is kept periodic. As a result, the density stratification, imposed body force, and interface-normal direction are aligned with x, while the sinusoidal perturbation is imposed along the periodic y-direction. The physics of the problem is unchanged; only the coordinate orientation is rotated.

2.4.3 Discontinuous Galerkin discretization and RK3

To simulate both the kinetic and PKPM models, we use the discontinuous Galerkin (DG) method. Within each cell, the solution is represented as a local polynomial expansion, and communication between neighboring cells is handled through numerical fluxes at shared interfaces [15]. For a conserved quantity $q(\mathbf{z}, t)$, the DG representation in a cell K maybe written as

$$q_h(\mathbf{z}, t)\big|_K = \sum_{j=1}^{N_p} q_j^K(t)\, \phi_j^K(\mathbf{z}), \tag{2.49}$$

where, ϕ_j^K are the local basis functions associated with the cell K and $q_j^K(t)$ are the corresponding time-dependent coefficients.

The DG method is well suited to the present problem because both the kinetic and PKPM equations contain strong advection-like structure and must be evolved in multidimensional phase or reduced phase space. In this work, both models use polynomial order $p = 2$ together with a third-order explicit Runge–Kutta time integrator. The two models therefore share the same underlying DG–RK time-advancement framework even though they solve different governing equations.

Both models are parallelized through configuration-space domain decomposition using MPI. At selected times, Gkeyll writes both global diagnostics and model-specific outputs needed for post-processing, including moment fields and integral quantities.

2.4.4 Basis functions and their role

The fully kinetic runs in this thesis use a serendipity basis, while the PKPM runs use a tensor-product basis. These are not different physical models; they are different polynomial representations within the same DG framework. The main consequence of the basis choice is a difference in the number of local degrees of freedom and therefore in computational cost.

In general, the tensor-product basis retains the full tensor-product set of polynomial modes in each element, while the serendipity basis removes some of the higher-order interior modes in order to reduce cost. For the same mesh, polynomial order, and adequate resolution, both bases approximate the same underlying solution. The practical difference is therefore primarily one of efficiency. In the present study, the serendipity basis is used for the higher-dimensional fully kinetic runs, while the tensor-product basis is used for the reduced PKPM system.

2.4.5 Velocity-space representation in the two models

Although both models operate within a DG framework, they do not discretize the same phase-space structure. In the fully kinetic model, the distribution function is defined on a grid in velocity space that corresponds to

$$(v_x, v_y) \in [-v_{x,\max}, v_{x,\max}] \times [-v_{y,\max}, v_{y,\max}], \tag{2.50}$$

with

$$N_{v_x} = N_{v_y} = 16. \tag{2.51}$$

This constitutes a 2x2v resolution of the neutral RTI problem. Although a more realistic higher-dimensional simulation domain would involve 2x3v, that is the same as the resolution in the fluid–kinetic RTI study of Rodman et al. [19], but such simulation would require far greater computational expense. We thus use the 2x2v Vlasov model for most our fully kinetic reference simulations and only carry out higher-dimensional calculations where more specifically called for.

PKPM is a bit more subtle. In its continuum form, it is a magnetization truncation of a full 3v problem, but it only explicitly discretizes 1 velocity component parallel to magnetic field, $v_\parallel$ [16]. The dynamics in those directions are kinetic, with $v_\parallel$ defined on the grid

$$v_\parallel \in [-v_{\max}, v_{\max}] \tag{2.52}$$

is explicitly gridded, with

$$N_{v_\parallel} = 16, \tag{2.53}$$

while the perpendicular dynamics are encoded through the reduced variables of the PKPM system rather than through a fully resolved perpendicular velocity mesh.

The core idea for this comparison is that the kinetic model resolves the full velocity phase-space grid, whereas the PKPM model has explicit kinetic behavior only in one dimension, with other directions accounted for through their integrated, moments of velocity. The calculation, therefore, is not between two identical sets of discrete grid points, but between a fully explicit description of the particle velocity space in kinetic form versus a partially implicit description using moments of particle velocity in kinetic form, both contained within the same DG framework.

2.4.6 Role of the computational framework in the present study

The shared computational framework provides a controlled setting for comparing the fully kinetic and PKPM models. Because both are implemented in Gkeyll, use DG discretization, evolve on the same configuration-space geometry, and use the same time-integration strategy, the comparison is less affected by unrelated numerical differences. This makes it easier to distinguish model-based differences from solver-based ones. The fully kinetic model provides access to the retained phase-space structure directly, while PKPM reduces the perpendicular degrees of freedom in order to lower computational cost. One purpose of the present work is to determine how much of the relevant RTI physics survives that reduction. The common computational framework is therefore an important part of the comparison itself.

2.5 Problem setup and initialization

2.5.1 Hydrostatic state

Both the fully kinetic and PKPM simulations are initialized from the same hydrostatic base state. The goal is to begin from a stratified density and pressure profile in mechanical balance under a constant imposed acceleration, and then to trigger Rayleigh–Taylor instability by adding a small perturbation. In the present implementation, the unstable direction is taken along x, so hydrostatic balance is written as

$$\frac{dp}{dx} = -nmg, \tag{2.54}$$

where n is the number density, m is the particle mass, and g is the magnitude of the imposed acceleration.

The density profile is prescribed using a hyperbolic tangent form,

$$n(x) = \frac{n_0}{2}\tanh\left(\frac{\alpha x}{L_x}\right) + \frac{3}{2}n_0, \tag{2.55}$$

where $n_0 = 0.5$, $\alpha = 25$, and L_x refers to the half-width of the grid domain in the stratified direction. his form provides a steep but smooth density change at the location of the

interface, transitioning from low to high density regions. The corresponding Atwood number is

$$\mathrm{At} = \frac{n_1 - n_2}{n_1 + n_2} = \frac{1}{3}, \tag{2.56}$$

which matches the value used in the reference kinetic RTI study of Rodman et al. [19].

The pressure profile is obtained by integrating the hydrostatic balance equation (Eq. (2.54)) using the prescribed density profile (Eq. (2.55)) as

$$p(x) = -gm\left[\frac{n_0}{2}\frac{L_x}{\alpha}\,\mathrm{lncosh}\left(\frac{\alpha x}{L_x}\right) + \frac{3}{2}n_0 x\right] + \frac{3}{2}n_0 T_0, \tag{2.57}$$

where T_0 must be chosen such that the pressure is always positive throughout the domain and for the given implementation we used the reference temperature is taken as

$$T_0 = \frac{2gm}{3\alpha}\left(3L_x\alpha + L_x\,\mathrm{lncosh}\,\alpha\right). \tag{2.58}$$

The resulting temperature profile is then obtained from the ideal relation

$$T(x) = \frac{p(x)}{n(x)}. \tag{2.59}$$

This base state satisfies hydrostatic balance, although, as discussed by Rodman et al., it is not a true collisional equilibrium distribution when collisions are present. At sufficiently low collisionality, the interface may therefore diffuse before the instability has time to develop [19].

2.5.2 Initialization of the fully kinetic model

Using the kinetic equation in the fully kinetic model involves preparing a particle distribution function at (t = 0) constructed from the hydrostatic base-state quantities of density, temperature and drift velocity. The fundamental process involves initializing the different particle species using local thermodynamic equilibrium distributions whose statistical moments match the values provided by the hydrostatic solution (Equations (2.55)–(2.59)).

In general form, this may be written as

$$f(\mathbf{v}; x, y, 0) = \frac{n(x)}{\left(2\pi v_{\mathrm{th}}^2(x)\right)^{D_v/2}} \exp\left[-\frac{|\mathbf{v} - \mathbf{u}(x, y, 0)|^2}{2v_{\mathrm{th}}^2(x)}\right], \tag{2.60}$$

where D_v represents the number of dimensions used in velocity space and

$$v_{\mathrm{th}}(x) = \sqrt{\frac{T(x)}{m}}. \tag{2.61}$$

In the present Vlasov implementation, this initialization is applied through Gkeyll's local thermodynamic equilibrium (LTE) projection so that the density and temperature are set from the hydrostatic profiles and the initial drift velocity carries the perturbation used to drive the instability. This is the same basic construction used in the neutral continuum-kinetic RTI study of Rodman et al., [19].

2.5.3 Initialization of the PKPM model

Initialization for the PKPM code comes from exactly the same density and temperature hydrostatic profiles but in reduced variables of the simplest (lowest order) model for a gyrotropic system. In this instance, the initial parallel kinetic component is set as

$$F_0(x, v_\parallel, 0) = \frac{n(x)}{\sqrt{2\pi T(x)}} \exp\left[-\frac{v_\parallel^2}{2T(x)}\right], \tag{2.62}$$

where normalized units with $m = 1$ have been used,

The second PKPM variable is initialized as

$$G(x, v_\parallel, 0) = T(x)\, F_0(x, v_\parallel, 0). \tag{2.63}$$

The quantity G is defined for the lowest-order PKPM to be equal to

$$G = \frac{T_\perp}{m}(F_0 - F_1), \tag{2.64}$$

Using the expression given in Eq (2.63) ensures that the perpendicular pressure is encapsulated within G without needing to solve a separate equation for $T_\perp$. The parallel kinetic component, expressed as shown in (2.62), represents an isotropic equilibrium, therefore, any agyrotropic contributions are eliminated in the lowest-order PKPM when beginning from such a state.

Thus, although the fully kinetic and PKPM models are initialized differently, both are constructed from the same hydrostatic density and temperature fields.

2.5.4 Perturbation

To initiate the instability in a controlled way, a small localized single-mode sinusoidal perturbation is imposed on the initial velocity profile. In the instability direction, x, the perturbation is applied to the x-component of the drift velocity and it varies in the periodic y direction as:

$$u_x(x, y, 0) = -0.1\, v_{\mathrm{th},c} \cos(ky) \exp\left(-\frac{x^2}{2x_r^2}\right), \tag{2.65}$$

where $v_{\mathrm{th},c}$ is a typical thermal speed at the interface center and

$$x_r = \frac{L_x}{10}. \tag{2.66}$$

The exponential profile is designed to isolate the disturbance near the interface and prevent the bulk domain from diverging drastically from the initially hydrostatically balanced state.

Here, the wavenumber k is chosen to be the total periodic length in the y direction for the single-mode perturbation, i.e.,

$$k = \frac{2\pi}{2L_y} = \frac{\pi}{L_y}. \tag{2.67}$$

The corresponding classical Rayleigh-Taylor growth time is

$$\tau_{\mathrm{RT}} = \frac{1}{\sqrt{k\,\mathrm{At}\,g}}, \tag{2.68}$$

and the simulations are generally run for a duration of

$$t_{\mathrm{end}} = 3\tau_{\mathrm{RT}}. \tag{2.69}$$

The perturbation amplitude is kept small enough that the initial stage of the evolution remains close to the intended background equilibrium. This makes it possible to compare the early growth behavior, where the instability remains in its linear phase before transitioning to nonlinear evolution, of the kinetic and PKPM models under matched conditions.

2.5.5 Boundary conditions and forcing

The boundary conditions are chosen so that the problem remains a single-mode Rayleigh–Taylor configuration with a localized interface and a periodic perturbation direction. Since

the instability develops along x, the boundaries at $x = x_{min}$ and $x = x_{max}$ use fixed-function reservoir-like conditions, in which the distribution is set at the boundary to a prescribed equilibrium state so that the domain is effectively coupled to an external reservoir [4], while the boundaries at $y = y_{min}$ and $y = y_{max}$ will be periodic.

The periodic y-direction ensures continuity of the imposed sinusoidal perturbation across the domain. In the x-direction, the reservoir-like boundary treatment maintains the stratified background state and helps prevent artificial boundary-driven evolution. Reflecting boundaries would introduce wall effects that are not part of the intended problem, while simple outflow boundaries would not preserve the hydrostatic profile near the domain edges. The fixed-function treatment is therefore used so that the instability develops primarily from the interface in the center of the domain rather than from the boundaries, as is expected and was done in the kinetic simulation in Rodman et al., which left the ghost cells at the boundary edges of the stratification constant with the hydrostatic profiles.

The instability arises from an applied, external acceleration along the negative x,

$$\mathbf{a} = (-g, 0, 0). \tag{2.70}$$

In the fully kinetic neutral-species simulations, this acceleration enters directly as a body force in the kinetic equation. In the PKPM model, the same imposed acceleration is included in the reduced kinetic–moment evolution. Although the two models evolve different equations, they are therefore driven by the same external forcing.

For the neutral fully kinetic runs, the field solve is disabled, leaving a body-force driven evolution together with the collision operator. In the PKPM calculations, a prescribed magnetic field is retained because the model depends on the decomposition of the dynamics into parallel and perpendicular directions relative to $\mathbf{B}$. In this work, that field is fixed in time, so it defines the PKPM closure direction without adding self-consistent electromagnetic evolution to the neutral RTI problem.

2.5.6 Collision model and collisionality parameterization

Instead of using the full Landau collision operator, the present work uses reduced velocity-space collision operators to represent collisional effects in a computationally tractable way.

For the matched kinetic–PKPM comparison, the most important common operator is the Lenard–Bernstein / Dougherty-type form, since it is available in both models within the Gkeyll framework [18, 6, 12, 16].

In shorthand form, this operator may be written as

$$C[f] = \nu \nabla_{\mathbf{v}} \cdot \left[(\mathbf{v} - \mathbf{u}) f + \frac{T}{m} \nabla_{\mathbf{v}} f \right], \tag{2.71}$$

where ν is the collision frequency, The operator combines a drag term with a velocity-space diffusion term while preserving the conservative structure needed for continuum kinetic and reduced-kinetic simulations.

In addition, a set of runs comparing kinetic-level runs are presented that use the BGK operator (Eq 2.25). These BGK runs are included to give some additional kinetic context for our cases, in cases where performing high-dimensional fully kinetic comparisons with BGK is far simpler than using the LBO operator. Nonetheless, the main comparison that is performed between PKPM and kinetic is done using the LBO operator.

Collisionality is measured using the Knudsen number, defined as:

$$\mathrm{Kn} = \frac{\lambda_m}{L_{\mathrm{ref}}}, \tag{2.72}$$

where λ_m is the mean free path and L_{ref} is the characteristic length scale used for the problem. The corresponding collision frequency is then chosen through

$$\nu = \frac{v_{\mathrm{th},c}}{\lambda_m} = \frac{v_{\mathrm{th},c}}{\mathrm{Kn}\, L_{\mathrm{ref}}}, \tag{2.73}$$

with $v_{\mathrm{th},c}$ derived from the characteristics of the reference thermal state in the plasma interface. Changing Kn thus provides a single, smooth control parameter to change collisionality, while holding the rest of the simulation parameters fixed.

The specific Knudsen numbers studied in the following chapters are chosen to compare more fluid-like and more transport-dominated regimes. At low Knudsen numbers, the kinetic distribution relaxes quickly to local equilibrium, such that behavior is much more fluid like. As the Knudsen number is increased, the mean free path becomes larger than L_{ref}, non-equilibrium effects such as kinetic energy and momentum transport, become more prominent and are critical to understanding evolution near the interface [19].

2.6 Diagnostics and post-processing

2.6.1 Overview of the analysis

The purpose of the post-processing pipeline is to convert the raw simulation outputs into a consistent set of diagnostics that can be compared across the fully kinetic and PKPM models. The analysis is not limited to visual inspection of contour plots. Instead, it is designed to extract physically meaningful measures of the instability, including interface evolution, fitted growth rate, thermal structure, heat flux, and the relative importance of ideal and non-ideal transport terms.

To accomplish this, we perform the following steps. Firstly, the DG output fields are reconstructed onto a plotting grid. Second, the density field is used to identify the evolving interface and extract a contour-based measure of the interface amplitude h(t). Third, a fitted growth rate is obtained from $ln(h(t))$ over a selected time interval. Finally, model-dependent moment fields are used to evaluate transport quantities and compare how the two descriptions represent the evolving RTI dynamics. This workflow follows the same general spirit as the diagnostic approach used in the kinetic RTI study of Rodman et al., while extending it to the kinetic–PKPM comparison carried out here [19].

2.6.2 Reconstruction of saved fields

The fields written by Gkeyll are stored in DG form, meaning that the saved data consist of local polynomial coefficients rather than pointwise values on a standard plotting grid. The first step in the post-processing is therefore to reconstruct the DG solution onto a uniform grid suitable for visualization and analysis. If $q_h(\mathbf{x}, t)$ denotes a saved DG field, then the reconstructed quantity used in the analysis is obtained through interpolation of the DG representation onto the plotting coordinates.

$$q_{ij}^{(n)} = \mathcal{I}_p[q_h]\big(x_i, y_j, t_n\big), \tag{2.74}$$

where $\mathcal{I}_p$ denotes the interpolation associated with the polynomial order p, (x_i, y_j) are the plotting-grid coordinates, and t_n is the time of the n-th saved frame.

The plotting coordinates are taken from the configuration-space grid associated with the output files. If $x_{i-1/2}$ and $x_{i+1/2}$ denote neighboring cell edges, then the cell-center location is given by

$$x_i = \frac{1}{2}\left(x_{i-1/2} + x_{i+1/2}\right), \qquad y_j = \frac{1}{2}\left(y_{j-1/2} + y_{j+1/2}\right). \tag{2.75}$$

All contour extraction, line profiles, and derived post-processing quantities are then evaluated from these reconstructed fields rather than directly from the modal coefficients.

2.6.3 Interface extraction from the density field

The main geometric diagnostic used for both models is the density interface. At each saved frame, the reconstructed density field is scanned row by row along the periodic y-direction, and the locations at which the density crosses a chosen contour value n^* are identified. For a fixed row y_j, the crossing locations satisfy

$$n_h(x, y_j, t_n) = n^*. \tag{2.76}$$

The crossings are obtained by interpolation between neighboring points that bracket the selected contour level. If a crossing lies between x_i and x_{i+1}, then

$$x_c = x_i + \frac{n^* - n_h(x_i, y_j, t_n)}{n_h(x_{i+1}, y_j, t_n) - n_h(x_i, y_j, t_n)}\left(x_{i+1} - x_i\right). \tag{2.77}$$

where n^* is chosen within the density range of the interface. The reference value $n^* = 0.6$ was selected from the density values encountered at the interface, so that the contour represents the same physical interface across different saved frames. The interface width is then determined from the extreme displacement of the interface in the unstable x-axis, defined as

$$x_{\min}(t_n) = \min x_c, \qquad x_{\max}(t_n) = \max x_c. \tag{2.78}$$

The resulting interface width or amplitude:

$$h_x(t_n) = x_{\max}(t_n) - x_{\min}(t_n). \tag{2.79}$$

This is equivalent to the bubble-spike measure introduced in the reference work [19, 22, 24, 23] , but is rotated here as the RTI evolves in the x-direction. The value is computed by

first extracting the set of vertices that comprise the contour from the reconstructed density field and then finding the maximum and minimum x coordinates in the set. A contour provides a more global measure of the overall bubble-and- spike perturbation and so would be preferred to make comparisons across different values of morphology or grid resolution.

2.6.4 Growth-rate estimation

To quantitatively measure the growth, one would typically analyze the time history of $h_x(t_n)$. In the roughly linear portion of the unstable region, the amplitude of the interface might be modeled as

$$h_x(t) \approx h_{x,0} e^{\gamma t}, \tag{2.80}$$

so that

$$\ln h_x(t) \approx \gamma t + C. \tag{2.81}$$

In the present work, the fit is carried out in normalized time using t/τ_{RT}, and only a selected fitting interval is retained. If the retained times are denoted by t_m, and

$$\ell_m = \ln h_x(t_m), \tag{2.82}$$

then the least-squares slope is computed from

$$\gamma^* = \frac{\sum_m (t_m - \bar{t})(\ell_m - \bar{\ell})}{\sum_m (t_m - \bar{t})^2}, \tag{2.83}$$

where $\bar{t}$ and $\bar{\ell}$ are averages over the fitting interval. The physical growth rate is then recovered as

$$\gamma = \frac{\gamma^*}{\tau_{\mathrm{RT}}}. \tag{2.84}$$

The scripts also compare the fitted growth rate to the ideal and theoretical references

$$\gamma_{\mathrm{ideal}} = \sqrt{kg\,\mathrm{At}}, \tag{2.85}$$

$$\gamma_{\mathrm{theory}} = \sqrt{kgA_t + \nu_v^2 k^4} - (\nu_v + \xi)k^2,$$

and report the ratio $\gamma/\gamma_{\mathrm{theory}}$. Since this study is concerned primarily with comparison against the expected kinetic or collisional behavior, this ratio is more informative than

comparison with the ideal inviscid growth rate alone. It directly shows how closely the measured instability growth follows the theoretical prediction for the regime of interest. Here, $\nu_v = v_{\mathrm{th},c}\lambda_m/2$ is the kinematic viscosity, and ξ represents the diffusive contribution included in the theory. As noted by Rodman et al., the earliest frames are typically excluded from the fit as the interface may not have fully transitioned to an exponential growth stage due to effects such as transient relaxation and collision diffusion [19].

2.6.5 Temperature-field diagnostics

The evolution of the thermal state is observed via saved LTE moments. We extract the relevant stored quantity T/m, once normalized, corresponds to the squared thermal-speed field,

$$v_{\mathrm{th}}^2 = \frac{T}{m}. \tag{2.86}$$

As LTE moment layout varies, the script is designed to identify the correct component, where we typically encounter it in component 4 for the (2x3v) layout and component 3 for the (2x2v) layout. We can then reconstruct this field on the same spatial grid as used for the density plots and iterate through it frame-by-frame, as depicted below. It is predominantly used as a background field to illustrate the behavior of the interface width and local thermal state during the instability.

2.6.6 Fully kinetic transport reconstruction

For kinetic (2x2v) results, we obtain transport terms by extrapolating (f) back out on the spatial phase grid: the relevant fields are recovered from $f(x_i, y_j, v_{x,\alpha}, v_{y,\beta}, t_n)$ at each saved frame. Since (f) are in DG format, they must be extrapolated back on the structured spatial and velocity grid using `postgkyl`.

The computation of the lowest moments done by numerical quadrature over the resolved velocity grid. Mathematically, these quantities are defined as velocity integrals, the density is obtained from

$$n(x_i, y_j, t_n) = \int f \, dv_x \, dv_y, \tag{2.87}$$

and the momentum densities are

$$m_x(x_i, y_j, t_n) = \int v_x f \, dv_x \, dv_y, \qquad m_y(x_i, y_j, t_n) = \int v_y f \, dv_x \, dv_y. \tag{2.88}$$

From these, the bulk velocities are reconstructed as

$$u_x = \frac{m_x}{n}, \qquad u_y = \frac{m_y}{n}. \tag{2.89}$$

These results represent the lowest moments calculated just before the formation of transport terms. Next, we translate the velocities to peculiar velocity by defining

$$w_x = v_x - u_x, \qquad w_y = v_y - u_y. \tag{2.90}$$

From peculiar velocities, we compute the gas-frame pressure tensor, components for pressure and stress:

$$P_{xx} = m \int w_x^2 f \, dv_x \, dv_y, \qquad P_{xy} = m \int w_x w_y f \, dv_x \, dv_y, \qquad P_{yy} = m \int w_y^2 f \, dv_x \, dv_y. \tag{2.91}$$

For $2x2v$ calculations, we take the two-dimensional analog of scalar pressure:

$$p = \frac{1}{2} \left(P_{xx} + P_{yy} \right), \tag{2.92}$$

and from this we then extract the non-ideal stresses:

$$\Pi_{xx} = P_{xx} - p, \qquad \Pi_{xy} = P_{xy}, \qquad \Pi_{yy} = P_{yy} - p. \tag{2.93}$$

So, at this stage, the script has already reconstructed the density, velocity, pressure, and non-ideal stress fields from the saved distribution.

By this stage, we have obtained density, velocity, pressure and the non-ideal stress terms. Next is the gas-frame heat flux calculated from the same reconstructed distribution function f, where we use the peculiar kinetic energy $w_x^2 + w_y^2$. The components are

$$q_x = \frac{m}{2} \int (w_x^2 + w_y^2) w_x f \, dv_x \, dv_y, \qquad q_y = \frac{m}{2} \int (w_x^2 + w_y^2) w_y f \, dv_x \, dv_y. \tag{2.94}$$

These quantities are the fully kinetic heat-flux contributions used in the transport plots. Finally, we construct the transport terms for plots: if the plots focus on the x direction, we calculate ideal enthalpy as

$$\mathcal{F}_x^{\text{enth}} = \frac{5}{2} u_x p, \tag{2.95}$$

the bulk kinetic-energy transport is

$$\mathcal{F}_x^{\text{bulk}} = \frac{1}{2} \rho u_x u^2, \tag{2.96}$$

the heat-flux contribution is

$$\mathcal{F}_x^{q} = q_x, \tag{2.97}$$

and the stress-work contribution is

$$\mathcal{F}_x^{\Pi} = u_x \Pi_{xx} + u_y \Pi_{xy}, \tag{2.98}$$

with $\rho = mn$ and $u^2 = u_x^2 + u_y^2$. The logic remains consistent when plotting the other direction. Thus, the script reconstructs the fields in sequence and then forms whichever transport term is requested for plotting.

2.6.7 PKPM reduced-moment diagnostics

For PKPM, post-processing relies on stored reduced-moment variables. In our setup, PKPM's results come from the two stored files using together: `pkpm_moms`, which stores reduced kinetic moments, and `pkpm_fluid`, which stores the bulk fluid quantities.

From `pkpm_fluid`, the bulk velocity components are reconstructed from

$$u_x = \frac{\rho u_x}{\rho}, \qquad u_y = \frac{\rho u_y}{\rho}, \qquad u_z = \frac{\rho u_z}{\rho}. \tag{2.99}$$

We use M_{ij} from the same `pkpm_fluid` to construct the clean pressure tensor

$$P_{ij} = M_{ij} - \rho u_i u_j. \tag{2.100}$$

The scalar pressure and non-ideal stress then follow from

$$p = \frac{1}{3}(P_{xx} + P_{yy} + P_{zz}), \qquad \Pi_{ij} = P_{ij} - p\delta_{ij}. \tag{2.101}$$

Using quantities from `pkpm_moms`, we read in the reduced ρ, $p_\parallel$, $p_\perp$, $q_\parallel$, $q_\perp$. We then construct the heat flux as

$$\mathbf{q} \approx (q_\parallel + q_\perp)\,\hat{\mathbf{b}}, \tag{2.102}$$

with $\hat{\mathbf{b}}$ being the unit vector along the fixed magnetic field. This is fully consistent with PKPM: in gyrotropic PKPM, the parallel and perpendicular heat flux components appear as field-parallel and field-perpendicular parts of the reduced heat flux [16].

Using the selected plotting direction $k \in x, y, z$, PKPM computes transport fields identical in form to the corresponding kinetic results

$$\mathcal{F}_k^{\mathrm{enth}} = \frac{5}{2} u_k p, \tag{2.103}$$

$$\mathcal{F}_k^{\mathrm{bulk}} = \frac{1}{2} \rho u_k u^2, \tag{2.104}$$

$$\mathcal{F}_k^{q} = q_k, \tag{2.105}$$

and

$$\mathcal{F}_k^{\Pi} = u_i \Pi_{ik}. \tag{2.106}$$

This produces PKPM transport plots of the same structure as those appearing as analogues to Figure 6 for the kinetic case of Rodman et al., while still accurately respecting the set of reduced variables evolved by PKPM.

2.6.8 Summary of the diagnostic strategy

In addition to the physical diagnostics, the runtime of each simulation is also recorded. This quantity is not used to define any physical observable; it is included only to compare the computational cost of the PKPM and fully kinetic models. Since one of the motivations for using PKPM is its potential to retain the essential transport physics at a much lower cost than a higher-fidelity kinetic simulation, runtime provides a practical measure of the benefit of model reduction. For this reason, it is discussed later together with the physical diagnostics, where it helps illustrate the accuracy-cost trade-off rather than serving as an independent diagnostic.

The diagnostics used in this thesis progress from simple geometric measures to transport-based quantities. The density field is used both to identify the interface and to define the

instability amplitude $h_x(t)$, from which the growth rate is extracted. In the fully kinetic runs, the LTE moments provide a concise way to track how the evolving plasma state affects the particle distribution, while the reconstructed $2x2v$ distribution is used to evaluate quantities such as pressure and heat flux. In PKPM, the `pkpm_moms` and `pkpm_fluid` outputs are used to reconstruct the reduced density, pressure, velocity moments, and the associated heat-flux and stress-work fields.

Chapter 3

RESULTS AND DISCUSSION

This chapter presents the numerical results from the fully kinetic and PKPM simulations and interprets them in terms of interface evolution, growth behavior, and transport effects. We begin by verifying that the two models are initialized consistently and evolve from equivalent physical conditions. This step is necessary because the later comparison is only meaningful if the observed differences can be attributed to the models themselves rather than to differences in setup, forcing, or boundary treatment.

3.1 Verification of the problem setup

Before comparing the fully kinetic and PKPM results, it is important to confirm that both models represent the same initial physical problem discussed in Chapter 2. In particular, we check that the density stratification, imposed perturbation, and thermal structure are initialized consistently, and that the early evolution reflects the intended setup rather than numerical or boundary-driven artifacts.

3.1.1 Consistency of the initialized state

Figure 3.1 shows that the fully kinetic and PKPM models are initialized from the same underlying physical state. In both cases, the density field has the same smooth hyperbolic-tangent transition across the stratified direction, centered near $x \approx 0$, with the lighter region at negative x and the denser region at positive x. This matches the intended initial configuration.

The perturbation velocity u_x also agrees closely between the two initializations. In both models, the perturbation is localized near the interface and varies sinusoidally along the periodic y-direction, as intended.

Temperature fields are again the same. In the fully kinetic case, the thermal quantity

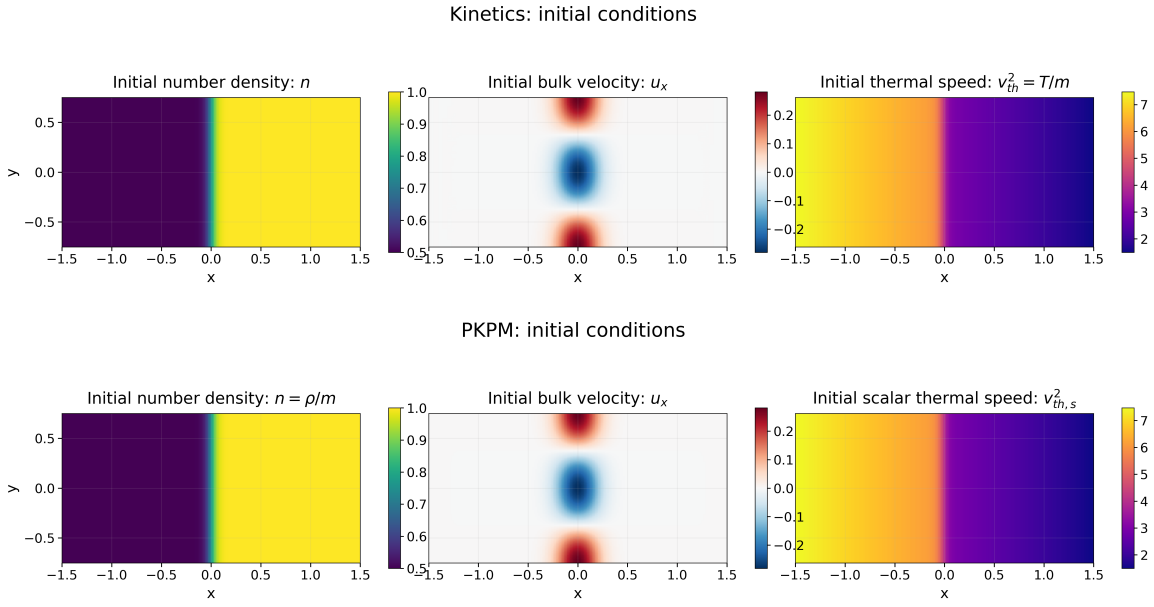

Figure 3.1: Initial-condition fields for the matched Kn = 0.01 case. Top: fully kinetic $2x2v$ LBO run. Bottom: PKPM LBO run with $\mathbf{B} \parallel \hat{\mathbf{x}}$. Shown are the density, the imposed perturbation velocity u_x, and a temperature-related field.

is shown as T/m, while in PKPM a scalar thermal field is constructed from the reduced pressures according to

$$v_{\mathrm{th,scalar}}^2 = \frac{p_{\parallel} + 2p_{\perp}}{3\rho}. \tag{3.1}$$

Although the two models use different evolved variables, the resulting thermal profiles are smooth and show the same structure across the interface and in the bulk regions.

The main purpose of Fig. 3.1 is therefore to confirm that both models begin from a consistent stratification, perturbation, and thermal state, so that the later comparison can be interpreted as a true model comparison.

3.1.2 Early evolution of the interface

Figure 3.2 shows the early evolution of the density field for the same case introduced in the previous figure. At $t/\tau_{\mathrm{RT}} = 0$, both runs begin from a nearly flat interface located around (x=0). Over time, the interface begins to deform in the right direction and the small initial

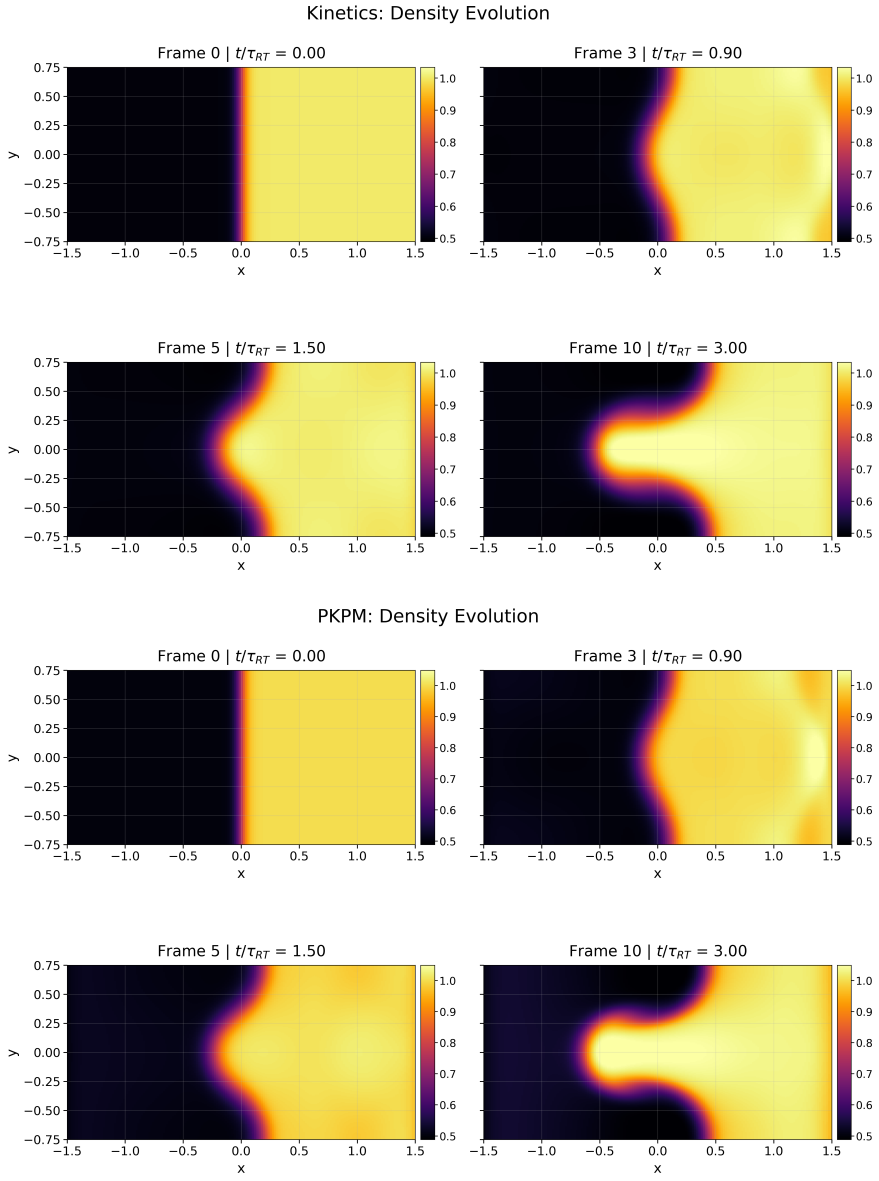

Figure 3.2: Early density evolution for the matched Kn = 0.01 case. Top: fully kinetic $2x2v$ LBO run. Bottom: PKPM LBO run with $\mathbf{B} \parallel \hat{\mathbf{x}}$. The panels show representative frames at $t/\tau_{\mathrm{RT}} = 0.00$, 0.90, 1.50, and 3.00.

perturbation grows so as to create a clearly resolved interfacial displacement.

By $t/\tau_{\mathrm{RT}} = 0.90$, the first differences appear between both runs: the interface has moved out along the middle of the domain ($y \approx 0$) while both the top and bottom sides still lag. This implies that the perturbation actually drove the instability. At $t/\tau_{\mathrm{RT}} = 1.50$, the shape of the interface becomes clearer and at $t/\tau_{\mathrm{RT}} = 3.00$, it is now developed and smooth. It should be emphasized that deformation only occurs around the interface, not out in the dense, uniform layers.

Additionally, the instability evolution appears symmetric about the middle plane because of the single-mode perturbation. Finally, we observe no spurious disturbance that arises at the domain boundary throughout these times. Moreover, we see similar early behavior in both models. The resulting shape of the interface is similar, as is the location of the greatest displacement, both models evolve from a nearly flat state into one with a larger, clearly developed displacement. Overall, this initial comparison suggests that both models successfully carry out the setup parameters, both numerically and conceptually, with no apparent inconsistencies or dominant spurious phenomena arising.

3.1.3 Verification of interface tracking

Figure 3.3 shows the interface-tracking contour method, which we used to determine the growth rate, behaves for the same two models. The white contour follows the density transition in a reasonable fashion throughout the entire evolution and continues to bound the minimum and maximum (x) values at either edge of the interface while the deformation grows. The interface is mostly planar at the beginning of the simulation ($t/\tau_{\mathrm{RT}} = 0$) and the corresponding amplitude is 0 in both cases, as is expected. The contour starts to move to the right near the middle of the grid at $t/\tau_{\mathrm{RT}} = 0.90$ and the two dashed lines capture this expansion.

At $t/\tau_{\mathrm{RT}} = 1.50$, the interface is more strongly deformed, but the contour remains smooth and single-valued over most of the domain, so the extracted $x_{\min}$ and $x_{\max}$ still provide a clear geometric measure of the instability amplitude. By $t/\tau_{\mathrm{RT}} = 3.00$, the shape of the interface has dramatically altered, nevertheless, the same contour is still faithfully

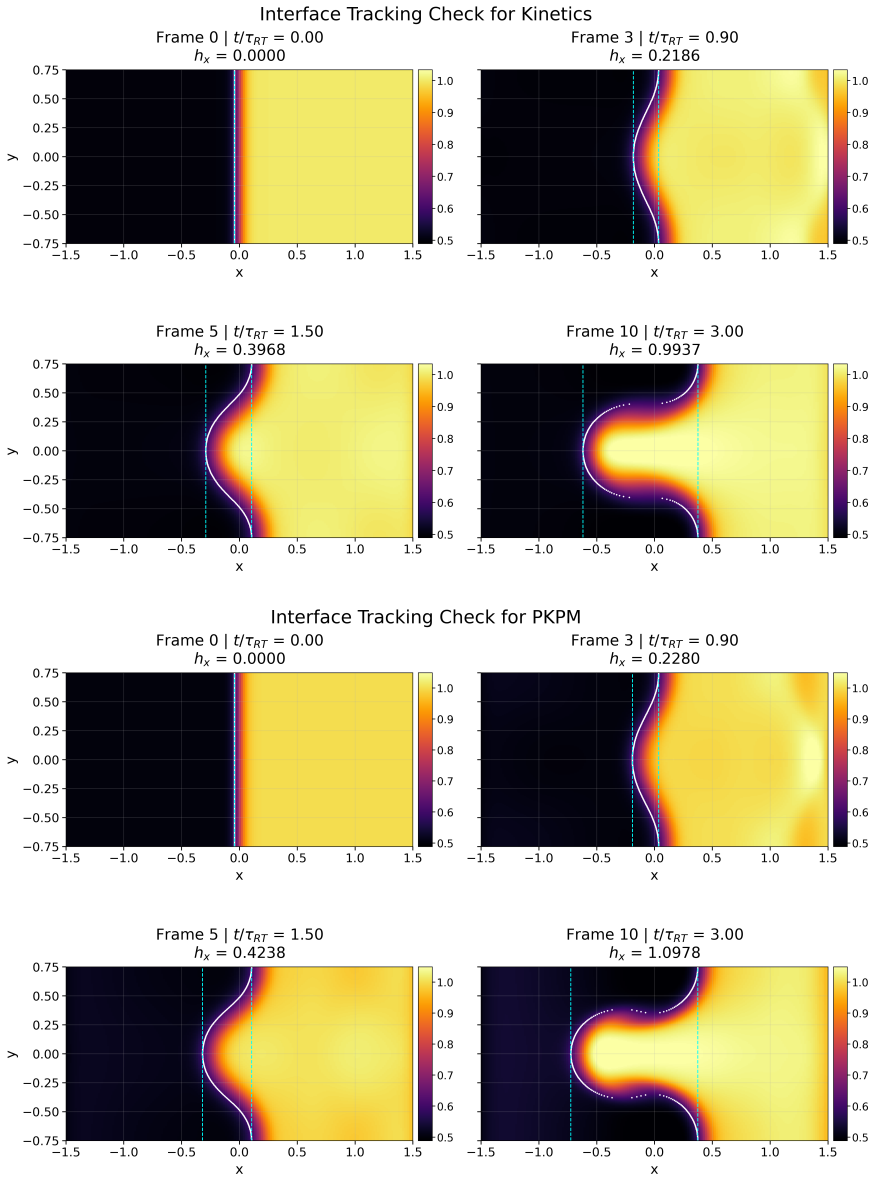

Figure 3.3: Contour-based interface tracking for the matched Kn = 0.01 case. Top: fully kinetic $2x2v$ LBO run. Bottom: PKPM LBO run with $\mathbf{B} \parallel \hat{\mathbf{x}}$. The white curve marks the selected iso-density contour used for interface tracking, while the dashed vertical lines indicate the extracted $x_{\min}$ and $x_{\max}$ positions used to define the interface amplitude $h_x = x_{\max} - x_{\min}$. Representative frames are shown at $t/\tau_{\mathrm{RT}} = 0.00$, 0.90, 1.50, and 3.00.

tracking the density boundary in both models, allowing the height measurement h_x to be taken in a consistently defined fashion.

Thus, the takeaway from Fig. 3.3 is the viability of using the given contour level of the density transition for analysis and that the derived quantities $x_{\min}$, $x_{\max}$, and h_x should serve as valid base quantities for analysis of the growth rate later on.

3.2 Comparison between the fully kinetic and PKPM results

Here we begin the first direct comparison between the fully kinetic reference run ($2x2v$, Kn = 0.01, LBO) and the reduced PKPM model at the same collisionality. For the PKPM case, we start with the B_{xy} orientation, defined by

$$\hat{\mathbf{b}} = \frac{\hat{\mathbf{x}} + \hat{\mathbf{y}}}{\sqrt{2}}.$$

It is helpful to consider the B_{xy} orientation because it represents the simplest form of the PKPM orientation that still maintains kinetic transport in both the x and y dimensions, which are the ones that matter most. Because the instability growth occurs in the x direction and variation along the y direction. For that reason, it is reasonable to expect B_{xy} to remain among the closer PKPM case to the fully kinetic reference, at least as an initial transport-based expectation. At the same time, that is only a starting hypothesis. The results later may show that different field orientations can agree more closely with the kinetic solution in specific quantities such as fitted growth rate or final interface morphology.

3.2.1 Density evolution and interface morphology

Figure 3.4 compares the density evolution of the fully kinetic and PKPM models for the matched Kn = 0.01 LBO case. At early times, the two solutions evolve in a similar manner: the initially planar interface deforms near the centerline, and the imposed perturbation develops into the expected single-mode structure with comparable symmetry and overall shape.

At later times, the same is retained in both models, but small differences emerge in the detailed interface geometry. Relative to the fully kinetic solution, the PKPM B_{xy} case shows a slightly different interface position and curvature near the head and shoulders, together

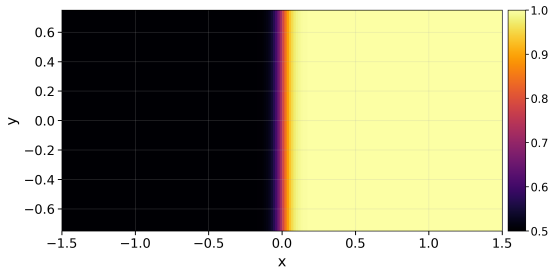

Kinetics, $t/\tau_{RT} = 0.00$

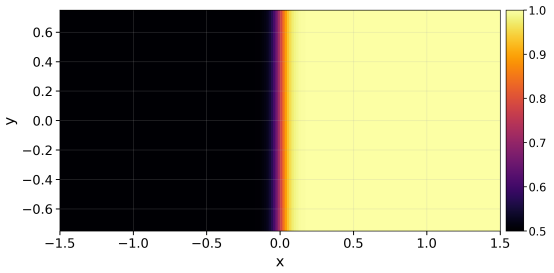

PKPM ($\mathbf{B} \parallel (\hat{\mathbf{x}} + \hat{\mathbf{y}})$), $t/\tau_{RT} = 0.00$

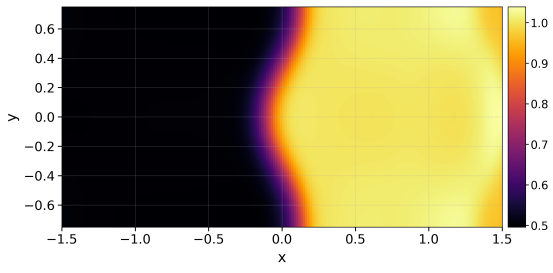

Kinetics, $t/\tau_{RT} = 0.90$

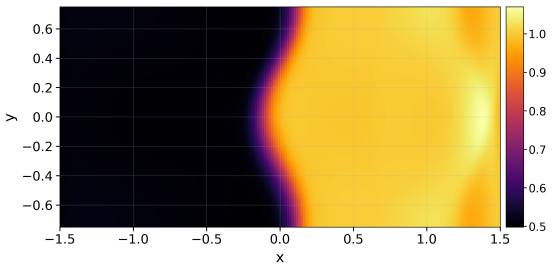

PKPM ($\mathbf{B} \parallel (\hat{\mathbf{x}} + \hat{\mathbf{y}})$), $t/\tau_{RT} = 0.90$

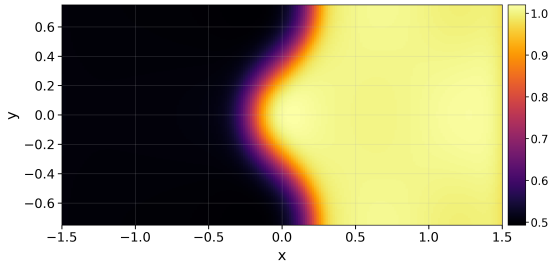

Kinetics, $t/\tau_{RT} = 1.50$

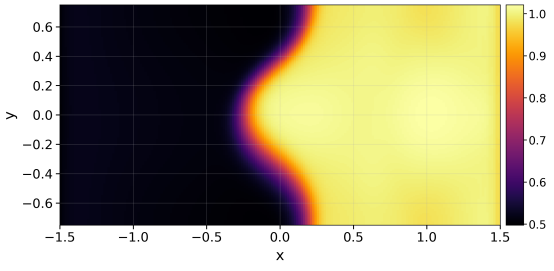

PKPM ($\mathbf{B} \parallel (\hat{\mathbf{x}} + \hat{\mathbf{y}})$), $t/\tau_{RT} = 1.50$

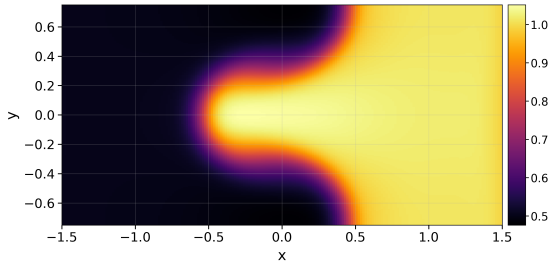

Kinetics, $t/\tau_{RT} = 3.00$

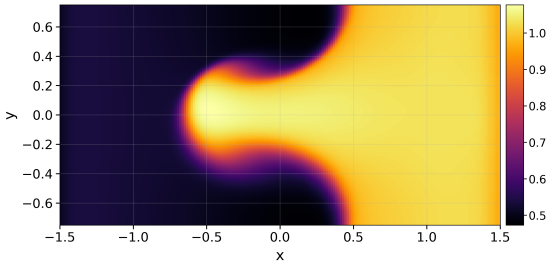

PKPM ($\mathbf{B} \parallel (\hat{\mathbf{x}} + \hat{\mathbf{y}})$), $t/\tau_{RT} = 3.00$

Figure 3.4: Comparison of density evolution between the fully kinetic reference and the PKPM B_{xy} case for Kn = 0.01 with LBO collisions. Each row compares the two models at the same normalized time. This layout makes it easier to assess differences in interface position, curvature, and transition-layer thickness at corresponding stages of the evolution.

with a modest difference in the thickness of the transition region. These differences remain limited in magnitude, but they mark the point at which the reduced description begins to depart from the fully kinetic reference. Overall, Fig. 3.4 shows that the PKPM B_{xy} case reproduces the main features of the fully kinetic interface evolution at Kn = 0.01.

3.2.2 Interface amplitude and growth-rate comparison

Figure 3.5 compares the interface amplitude and fitted growth rate for the fully kinetic model and the PKPM B_{xy} case. The two solutions remain close throughout the run, but the separation between them grows gradually with time. By the final frame, the interface amplitude reaches $h_x = 0.9937$ in the fully kinetic case and $h_x = 1.1064$ in PKPM, indicating that the PKPM front has advanced slightly farther in the unstable direction.

The fitted growth rates show the same trend. The fully kinetic case gives $\gamma = 0.7779$, while the PKPM B_{xy} case gives $\gamma = 0.7887$, so the PKPM growth is only slightly larger. This difference is small, but it is consistent with the density and interface-tracking plots, where the PKPM front remains modestly ahead at later times. At the same time, the two values are close enough to show that the reduced model is capturing the main growth behavior of the fully kinetic reference reasonably well.

Table 3.2 shows that the fitted growth rates for the two models are very close to one another and both remain below the theoretical reference value suggesting the evolution in both models is being moderated by effects that are not present in the ideal estimate, such as diffusion. The PKPM B_{xy} case is slightly closer to γ_{theory} than the fully kinetic case, while the high R^2 values in both fits indicate that the chosen fitting window is well described by an approximately linear trend in $\log(h_x)$.

3.2.3 Thermal-field evolution

We compare the thermal-field time evolution for the fully kinetic reference and the PKPM B_{xy} case for selected times, namely $t/\tau_{\text{RT}} = 0.00$, 0.90, 1.50 and 3.00. We plot T/m in the kinetic run and the scalar $(p_{\parallel} + 2p_{\perp})/(3\rho)$ in PKPM. Although these variables differ, they play equivalent roles as indicators of local temperature.

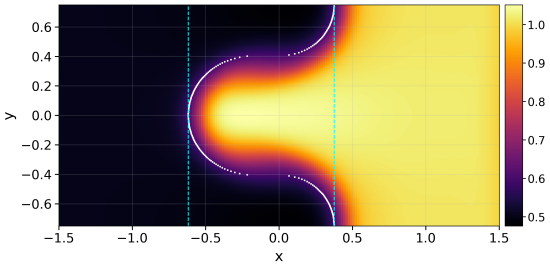

(a) Fully kinetic $2x2v$ LBO case at $t/\tau_{\mathrm{RT}} = 3.00$.

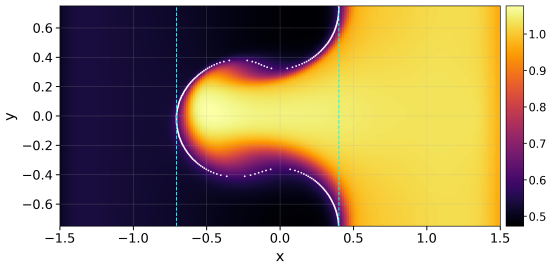

(b) PKPM LBO case with $\mathbf{B} \parallel (\hat{\mathbf{x}} + \hat{\mathbf{y}})$ at $t/\tau_{\mathrm{RT}} = 3.00$.

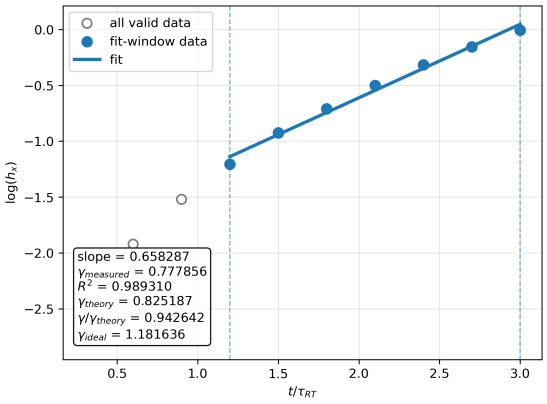

(c) Growth-rate fit for the fully kinetic case.

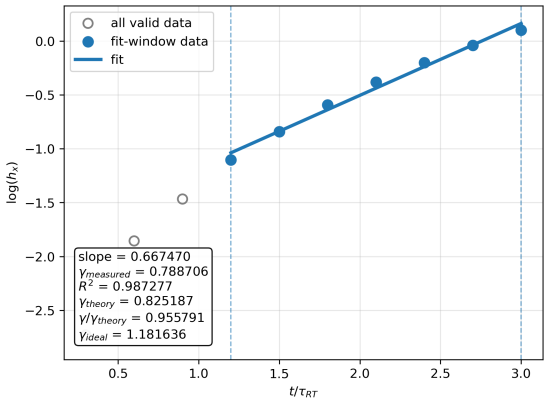

(d) Growth-rate fit for the PKPM B_{xy} case.

Figure 3.5: Comparison of interface amplitude growth between the fully kinetic model and the PKPM B_{xy} case for the matched $\mathrm{Kn} = 0.01$ LBO runs. The top row shows the tracked interface at the final saved time, $h_x = x_{\max} - x_{\min}$. The bottom row shows the corresponding logarithmic growth fits over the selected fitting interval.

Table 3.1: Frame-by-frame comparison of the interface amplitude h_x for the fully kinetic model and the PKPM B_{xy} case for the matched Kn = 0.01 LBO runs. The last column gives the difference $\Delta h_x = h_{x,\mathrm{PKPM}} - h_{x,\mathrm{kinetic}}$.

t/τ_{RT}	$h_{x,\mathrm{kinetic}}$	$h_{x,\mathrm{PKPM}}$	Δh_x
0.0	0.0000	0.0000	0.0000
0.3	0.0649	0.0652	0.0003
0.6	0.1463	0.1563	0.0100
0.9	0.2186	0.2306	0.0120
1.2	0.2995	0.3319	0.0324
1.5	0.3968	0.4312	0.0344
1.8	0.4921	0.5526	0.0605
2.1	0.6059	0.6825	0.0766
2.4	0.7283	0.8180	0.0897
2.7	0.8566	0.9608	0.1042
3.0	0.9937	1.1064	0.1127

Table 3.2: Summary of growth-rate quantities for the fully kinetic reference and the PKPM B_{xy} case for the matched Kn = 0.01 LBO runs.

Case	$\gamma_{\mathrm{measured}}$	γ_{theory}	$\gamma/\gamma_{\mathrm{theory}}$	R^2
Fully kinetic $2x2v$	0.7779	0.8252	0.9426	0.9893
PKPM B_{xy}	0.7887	0.8252	0.9558	0.9873

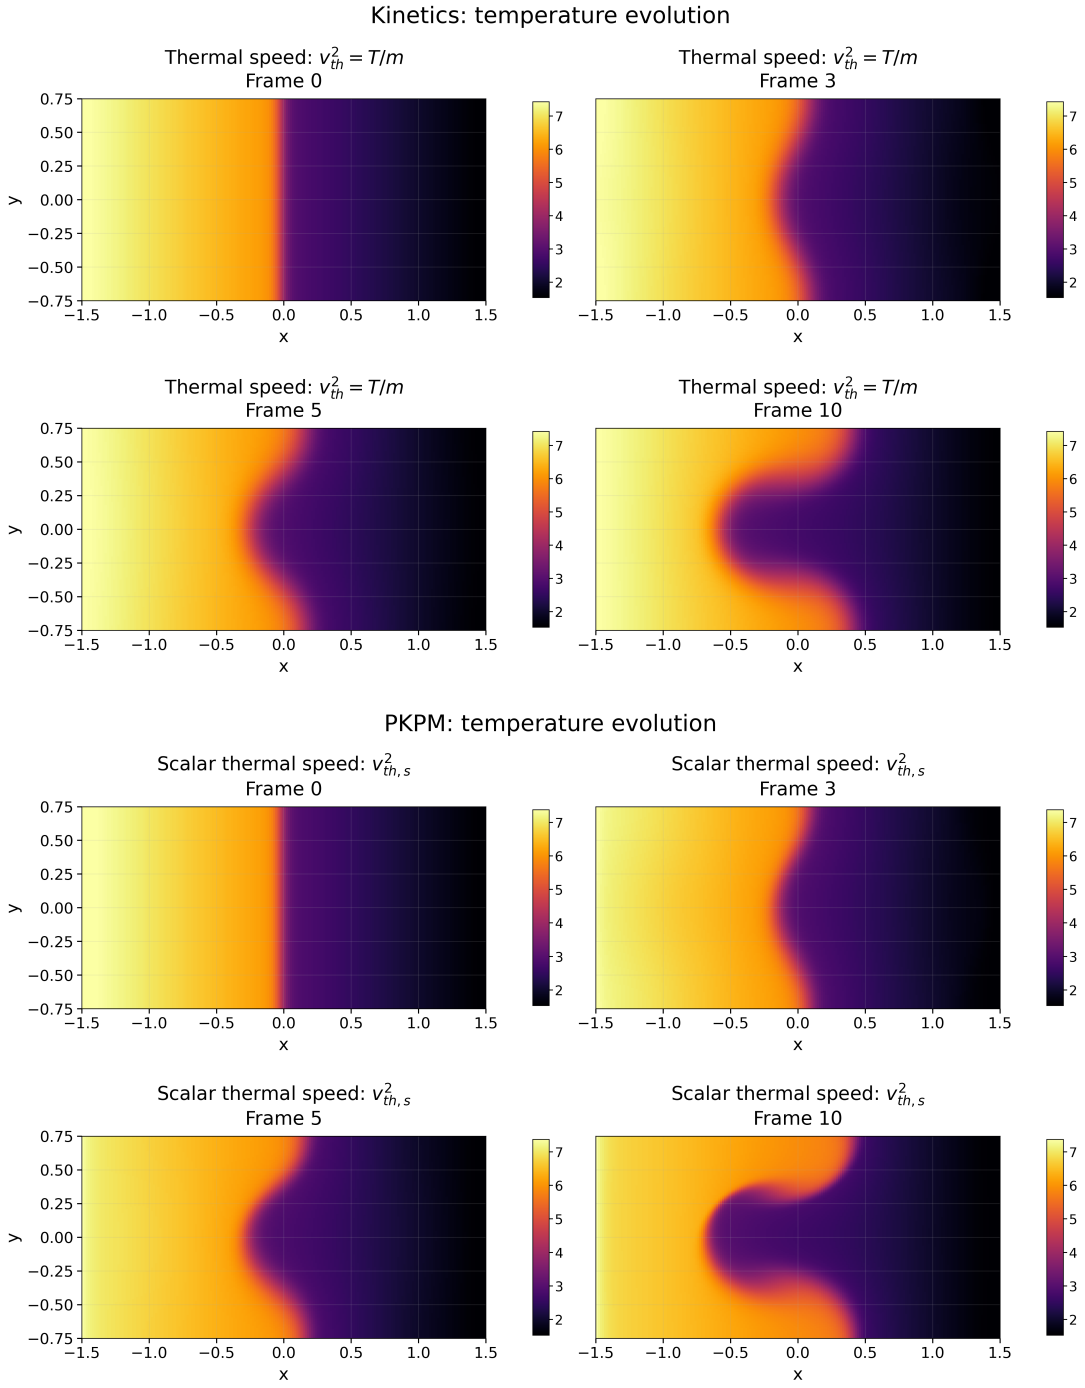

Figure 3.6: Thermal-field evolution for the matched Kn = 0.01 LBO comparison. Top: fully kinetic $2x2v$ case, plotted as T/m. Bottom: PKPM B_{xy} case, plotted using the scalar thermal quantity $(p_\parallel + 2p_\perp)/(3\rho)$. Representative frames are shown at $t/\tau_{\mathrm{RT}} = 0.00$, 0.90, 1.50, and 3.00.

As we see in the density evolution (Fig. 3.6), the temperature distribution has the same overall spatial features as density just like in the kinetic RTI result of Rodman et al., The high-temperature region remains on the low-density (right) side, while the low-temperature region sweeps along with the dense front as the interface proceeds toward negative x. The two models remain close throughout the selected frames. At $t/\tau_{\mathrm{RT}} = 0$, both show the same sharp thermal transition across the interface. By $t/\tau_{\mathrm{RT}} = 0.90$ and 1.50, the thermal front bends with the interface in the same way as the density field, and by $t/\tau_{\mathrm{RT}} = 3.00$ both runs show a similar cooler protrusion extending toward negative x. The differences are the PKPM front is slightly different in curvature and transition thickness near the advancing head.

3.2.4 Energy-flux decomposition for both the PKPM B_{xy} and Kinetic cases

The x-directed energy flux was decomposed into ideal and non-ideal contributions as

$$F_{E,x} = \underbrace{\frac{5}{2}u_x p}_{I_1} + \underbrace{\frac{1}{2}\rho u_x u^2}_{I_2} + \underbrace{q_x}_{\text{heat flux}} + \underbrace{u_i \Pi_{ix}}_{\text{stress-work}} , \tag{3.2}$$

where

$$F_{E,x}^{\text{ideal}} = I_1 + I_2, \qquad F_{E,x}^{\text{nonideal}} = q_x + u_i \Pi_{ix}. \tag{3.3}$$

The volume-averaged magnitudes of each contribution over the full computational domain are summarized in Table 3.3. The ratio reported in the final column is

$$R = \frac{\left\langle |F_{E,x}^{\text{ideal}}| \right\rangle}{\left\langle |F_{E,x}^{\text{nonideal}}| \right\rangle}. \tag{3.4}$$

Table 3.3: Whole-domain averaged energy-flux terms for both Kinetics and PKPM B_{xy} case at frame 10.

Model	$\langle\lvert I_1\rvert\rangle$	$\langle\lvert I_2\rvert\rangle$	$\langle\lvert q_x\rvert\rangle$	$\langle\lvert u_i\Pi_{ix}\rvert\rangle$	$\langle\lvert\text{ideal}\rvert\rangle$	$\langle\lvert\text{nonideal}\rvert\rangle$	R
Kinetics	7.10472×10^{-2}	3.51157×10^{-4}	8.94942×10^{-4}	4.24415×10^{-5}	7.13983×10^{-2}	9.37384×10^{-4}	76.16765
PKPM	8.63811×10^{-2}	6.76262×10^{-4}	1.28593×10^{-3}	1.72502×10^{-5}	8.70574×10^{-2}	1.30318×10^{-3}	66.80369

The dominant contribution to the ideal energy flux in both models is the pressure/enthalpy transport term,

$$I_1 = \frac{5}{2} u_x p. \tag{3.5}$$

For the kinetic simulation, this term contributes approximately

$$\frac{\langle |I_1| \rangle}{\langle |\text{ideal}| \rangle} = \frac{7.10472 \times 10^{-2}}{7.13983 \times 10^{-2}} \approx 99.51\%, \tag{3.6}$$

while for PKPM,

$$\frac{\langle |I_1| \rangle}{\langle |\text{ideal}| \rangle} = \frac{8.63811 \times 10^{-2}}{8.70574 \times 10^{-2}} \approx 99.22\%. \tag{3.7}$$

Therefore, the ideal energy transport is almost entirely controlled by the pressure/enthalpy flux in both simulations. The kinetic-energy flux term,

$$I_2 = \frac{1}{2} \rho u_x u^2, \tag{3.8}$$

is larger in PKPM than in the kinetic case, but it remains small compared with I_1. In the kinetic simulation, I_2 contributes only about 0.49% of the ideal flux, while in PKPM it contributes about 0.78%. Thus, even though I_2 is nearly doubled in PKPM, it does not control the overall ideal energy-flux behavior.

The ideal flux itself is stronger in PKPM:

$$\frac{\langle |\text{ideal}| \rangle_{\text{PKPM}} - \langle |\text{ideal}| \rangle_{\text{kinetic}}}{\langle |\text{ideal}| \rangle_{\text{kinetic}}} = \frac{8.70574 \times 10^{-2} - 7.13983 \times 10^{-2}}{7.13983 \times 10^{-2}} \approx 21.93\%. \tag{3.9}$$

Therefore, the smaller ideal-to-nonideal ratio in PKPM is not caused by a weaker ideal flux. Instead, PKPM has a larger ideal flux than the kinetic simulation.

The main difference comes from the non-ideal flux. The PKPM non-ideal contribution is

$$\frac{\langle |\text{nonideal}| \rangle_{\text{PKPM}} - \langle |\text{nonideal}| \rangle_{\text{kinetic}}}{\langle |\text{nonideal}| \rangle_{\text{kinetic}}} = \frac{1.30318 \times 10^{-3} - 9.37384 \times 10^{-4}}{9.37384 \times 10^{-4}} \approx 39.02\%, \tag{3.10}$$

which is a stronger relative increase than the ideal flux. This explains why the ideal-to-nonideal ratio decreases from

$$R_{\text{kinetic}} = 76.17 \tag{3.11}$$

to

$$R_{\mathrm{PKPM}} = 66.80. \tag{3.12}$$

Equivalently,

$$\frac{R_{\mathrm{PKPM}}}{R_{\mathrm{kinetic}}} = \frac{66.80369}{76.16765} \approx 0.877, \tag{3.13}$$

meaning that the PKPM case is about 12.3% less ideal-dominated than the kinetic case.

The non-ideal flux is mainly controlled by the heat-flux term q_x. In the kinetic simulation,

$$\frac{\langle |q_x| \rangle}{\langle |\mathrm{nonideal}| \rangle} = \frac{8.94942 \times 10^{-4}}{9.37384 \times 10^{-4}} \approx 95.47\%, \tag{3.14}$$

while in PKPM,

$$\frac{\langle |q_x| \rangle}{\langle |\mathrm{nonideal}| \rangle} = \frac{1.28593 \times 10^{-3}}{1.30318 \times 10^{-3}} \approx 98.68\%. \tag{3.15}$$

Thus, the non-ideal energy flux in both simulations is primarily a heat-flux contribution, but this dominance is even stronger in PKPM. The PKPM heat flux is approximately

$$\frac{\langle |q_x| \rangle_{\mathrm{PKPM}} - \langle |q_x| \rangle_{\mathrm{kinetic}}}{\langle |q_x| \rangle_{\mathrm{kinetic}}} = \frac{1.28593 \times 10^{-3} - 8.94942 \times 10^{-4}}{8.94942 \times 10^{-4}} \approx 43.69\% \tag{3.16}$$

larger than the kinetic heat flux. This increase is the primary reason why the PKPM ideal-to-nonideal ratio is smaller.

The stress-work term behaves differently. In the kinetic simulation,

$$\langle |u_i \Pi_{ix}| \rangle_{\mathrm{kinetic}} = 4.24415 \times 10^{-5}, \tag{3.17}$$

whereas in PKPM,

$$\langle |u_i \Pi_{ix}| \rangle_{\mathrm{PKPM}} = 1.72502 \times 10^{-5}. \tag{3.18}$$

Therefore, the PKPM stress-work contribution is approximately

$$\frac{1.72502 \times 10^{-5} - 4.24415 \times 10^{-5}}{4.24415 \times 10^{-5}} \approx -59.36\% \tag{3.19}$$

smaller than the kinetic value. This indicates that the reduction in the PKPM ratio is not caused by stronger stress-work. In fact, the opposite is observed: PKPM has a weaker stress-work contribution, while its heat-flux contribution is larger.

The spatial structure of the heat-flux plots supports this interpretation. In both simulations, q_x is concentrated near the deformed Rayleigh–Taylor interface rather than being

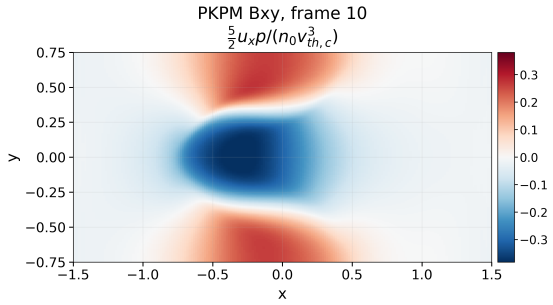

(a) PKPM: $\frac{5}{2}u_x p/(n_0 v_{th,c}^3)$

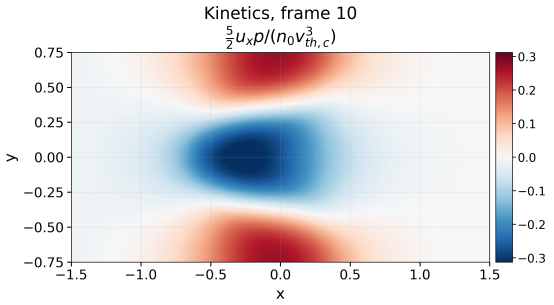

(b) Kinetics: $\frac{5}{2}u_x p/(n_0 v_{th,c}^3)$

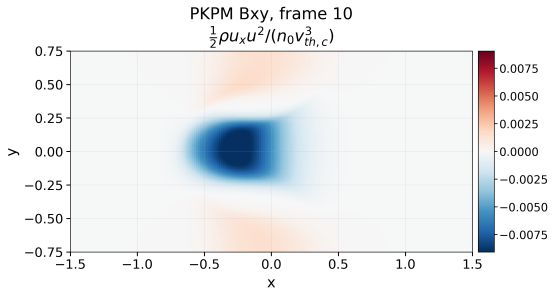

(c) PKPM: $\frac{1}{2}\rho u_x u^2/(n_0 v_{th,c}^3)$

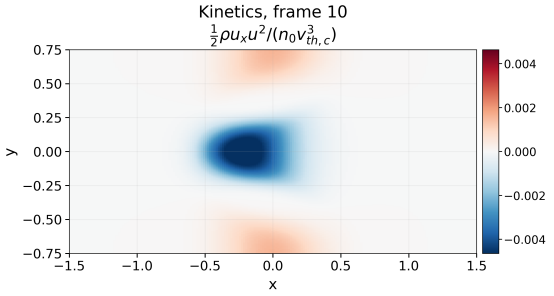

(d) Kinetics: $\frac{1}{2}\rho u_x u^2/(n_0 v_{th,c}^3)$

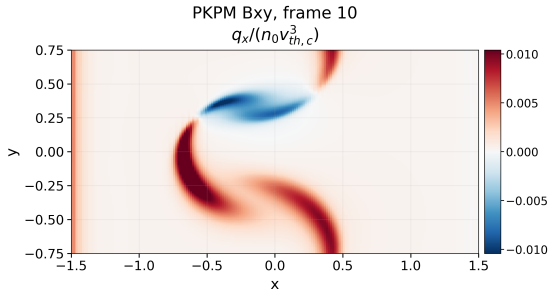

(e) PKPM: $q_x/(n_0 v_{th,c}^3)$

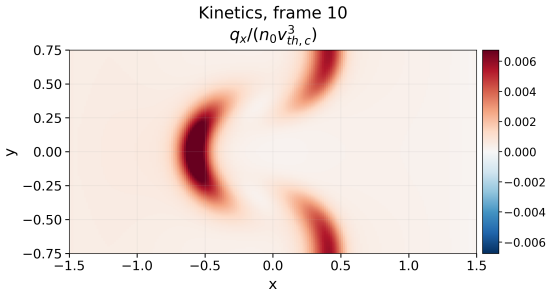

(f) Kinetics: $q_x/(n_0 v_{th,c}^3)$

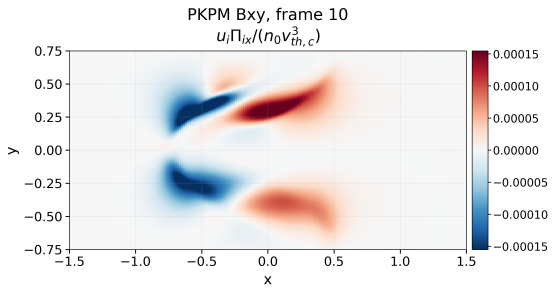

(g) PKPM: $u_i \Pi_{ix}/(n_0 v_{th,c}^3)$

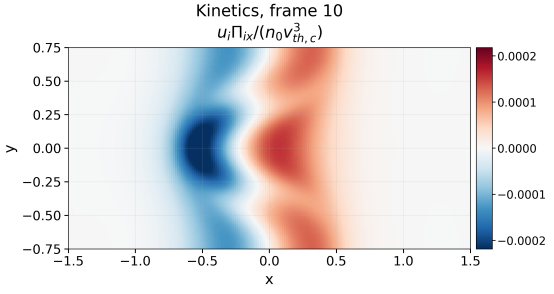

(h) Kinetics: $u_i \Pi_{ix}/(n_0 v_{th,c}^3)$

Figure 3.7: Side-by-side comparison of the ideal and non-ideal x-directed energy-flux terms for the fully kinetic and PKPM B_{xy} cases at frame 10. The top two rows show the ideal contributions, $\frac{5}{2}u_x p$ and $\frac{1}{2}\rho u_x u^2$, and the bottom two rows show the non-ideal contributions, q_x and $u_i \Pi_{ix}$.

distributed uniformly across the domain. This indicates that heat flux is generated mainly in regions where the density, pressure, and temperature fields are strongly distorted by the instability. In the PKPM case, the q_x field shows stronger alternating red and blue regions near the interface. These alternating structures suggest stronger local redistribution of thermal energy around the mixing layer. The kinetic heat-flux field shows the same general localization around the interface, but the structure is smoother and has a smaller mean magnitude.

This behavior can be interpreted as the heat flux is strongly coupled to the evolving interface. As the Rayleigh-Taylor instability deforms the density layer, sharper thermal and pressure gradients develop near the bubble and spike regions. These gradients produce localized heat-flux structures. In PKPM, the stronger heat-flux response may contribute to enhanced thermal redistribution near the interface, which can influence the detailed morphology of the mixed layer.

The stress-work plots show a different behavior. The kinetic stress-work field is broader and more coherent, while the PKPM stress-work field is more localized and patchy, with smaller positive and negative regions near the interface. This suggests that the PKPM stress-work contribution is less coherent and less dynamically important for the total non-ideal energy flux. Since the stress-work term contributes only about 4.53% of the kinetic non-ideal flux and only about 1.32% of the PKPM non-ideal flux, it is not the controlling term in the energy-flux balance.

Overall, both cases at frame 10 remains strongly ideal-flux dominated in both simulations. However, PKPM is relatively less ideal-dominated than the kinetic model. This is because the PKPM non-ideal flux increases more strongly than the PKPM ideal flux. More specifically, the ideal flux increases by about 21.9%, while the non-ideal flux increases by about 39.0%. The larger PKPM non-ideal flux is almost entirely due to the heat-flux term q_x, not the stress-work term.

3.3 Effect of magnetic-field orientation in PKPM

Before comparing different magnetic-field directions in PKPM, it is useful to clarify what **B** represents in the present setup. Here, the field is used primarily to define the direction

along which PKPM retains the kinetic dynamics, while the perpendicular response is treated through reduced moments.

When the field direction is fixed and only the magnitude is varied, for example from $B = 0.1$ to $B = 10$, the density evolution and fitted growth rates remain unchanged. In other words, the magnetic-field magnitude is not controlling the instability dynamics in these neutral species runs. This also helps the comparison with the fully kinetic reference since the kinetic runs do not include an active magnetic field. The results in this section show how the PKPM solution changes when the direction of the retained kinetic response is rotated relative to the interface.

3.3.1 Growth-rate and morphology trends across field orientations

After confirming that changing the magnitude of $\mathbf{B}$ does not affect the PKPM solution, the meaningful comparison is between different field directions. In the present neutral setup, $\mathbf{B}$ does not act through magnetic tension or Lorentz forcing. Instead, its role is to define the field-aligned kinetic direction in PKPM and therefore to determine how the RTI dynamics are divided between parallel kinetic evolution and perpendicular moment closure.

Table 3.4 summarizes the growth measurements for the kinetic and PKPM cases. The kinetic $2x2v$ LBO run remains the fairest direct reference for PKPM because both models use the same collision operator. The kinetic $2x2v$ BGK run is also useful because it shows how much the instability can change when only the collision model is altered. In the present results, the BGK case gives both a smaller fitted growth rate and a smaller final interface height than the kinetic LBO case. Its interface head remains narrower and less expanded at late time. A plausible reason is that the BGK operator relaxes the distribution more directly toward a local Maxwellian, which tends to suppress velocity-space structure and smooth the interface more strongly, whereas the LBO operator redistributes the distribution through drag and diffusion in velocity space and allows more detailed non-equilibrium structure to persist. In these runs, that difference appears as a weaker overall interface advance in BGK than in LBO. [3, 6, 12]

Table 3.4: Summary of kinetic and PKPM interface-growth measurements for Kn = 0.01.

Case	$h_x(t/\tau_{\mathrm{RT}} = 3)$	γ	$\gamma/\gamma_{\mathrm{ideal}}$	R^2
Kinetic $2x2v$, LBO	0.9937	0.7779	0.9426	0.9893
Kinetic $2x2v$, BGK	0.7599	0.7005	0.8489	0.9936
PKPM B_x	1.0978	0.7893	0.9565	0.9897
PKPM B_y	1.0691	0.7936	0.9617	0.9892
PKPM B_z	1.1655	0.8038	0.9741	0.9860
PKPM B_{xy}	1.1064	0.7887	0.9558	0.9873
PKPM B_{yz}	1.1329	0.8014	0.9712	0.9871
PKPM B_{xz}	1.1376	0.7922	0.9601	0.9877
PKPM B_{xyz}	1.1308	0.7935	0.9616	0.9868

The measured growth-rate ordering is

$$B_z > B_{yz} > B_y \approx B_{xyz} > B_{xz} > B_x > B_{xy} > \text{kinetic LBO} > \text{kinetic BGK},$$

while the final interface-height ordering is

$$B_z > B_{xz} > B_{yz} > B_{xyz} > B_{xy} > B_x > B_y > \text{kinetic LBO} > \text{kinetic BGK}.$$

So all PKPM cases grow faster and reach larger final amplitudes than both kinetic references, but they do not behave identically among themselves. The in-plane cases B_x and B_{xy} remain the closest to the kinetic LBO reference in fitted growth rate, whereas the cases containing a z-component move farther away. In the pure B_z case, the retained kinetic direction is not absent, but it lies out of the resolved x–y plane. As a result, the in-plane RTI dynamics are not moderated by an in-plane kinetic transport channel in the same way as in the B_x, B_y, or B_{xy} cases. Instead, more of the resolved in-plane evolution is handled through the reduced perpendicular response. In that sense, the B_z case behaves more fluid-like in the plane of the instability, which helps explain why it departs most strongly from the kinetic reference and develops the broadest, most mushroom-like interface head, giving both the largest γ and the largest final interface height.

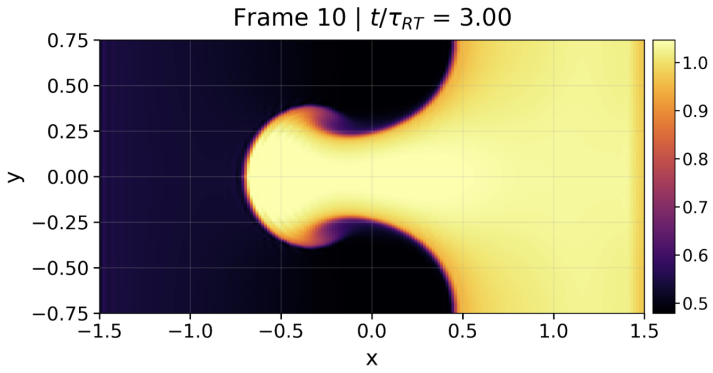

Figure 3.8: Density field for the PKPM B_z case at frame 10 ($t/\tau_{\mathrm{RT}} = 3.0$). Relative to the in-plane field-direction cases, the interface head is broader and more bulbous, giving the front a more pronounced mushroom-like shape. This case also corresponds to the largest fitted growth rate and final interface height among the PKPM orientations considered.

This same ordering is reflected in the interface morphology. The in-plane cases remain more compact and stay closer to the kinetic shape, especially near the advancing head and the upper and lower shoulders. Once a z-component is introduced, the head becomes broader and the late-time front develops a more mushroom-like profile. This is strongest in the pure B_z case (Fig. 3.8), and it remains visible, though less strongly, in the mixed B_{yz}, B_{xz}, and B_{xyz} cases. The out-of-plane field directions therefore do not just increase the fitted growth rate; they also change the nonlinear shape of the front.

The B_y case is a useful detail. Its fitted growth rate is not the smallest among the PKPM runs, yet it gives the smallest final interface height in the PKPM set. This indicates that the fitted linear growth rate is not the only quantity controlling the late-time amplitude. When the field-aligned kinetic direction is along the perturbation direction, the interface head appears to remain more compact and undergo less late-time widening. In other words, the B_y orientation does not strongly reduce the fitted exponential growth over the chosen window, but it does limit the later nonlinear swelling of the front. That is why its final h_x is smaller than in the other PKPM cases even though its fitted γ is not the smallest.

The overall trend is consistent with the structure of the PKPM model. When $\mathbf{B}$ lies in the x–y plane, part of the active RTI motion is still represented through the field-aligned

Table 3.5: Interpretation of kinetic and PKPM cases for the neutral $q = 0$ simulations.

Case	Interpretation
Kinetic LBO	Primary comparison case for PKPM. Higher growth than kinetic BGK and a more advanced interface, consistent with stronger retained non-equilibrium structure.
Kinetic BGK	Lowest-growth case in the set. The direct relaxation toward a local Maxwellian gives a narrower head and smaller final interface advance than the LBO case.
B_x	In-plane field aligned with interface growth. One of the closest PKPM cases to the kinetic LBO reference.
B_y	In-plane field aligned with the perturbation direction. Smallest final PKPM interface height, suggesting reduced late-time widening even though its fitted growth rate is not the smallest.
B_z	Field fully out of plane. Gives the largest growth rate, largest final interface height, and the most pronounced mushroom-like head.
B_{xy}	In-plane mixed direction. Closest PKPM case to the kinetic LBO reference in fitted growth.
B_{yz}	Mixed case with out-of-plane component. Stronger growth and broader head than the in-plane cases, but less extreme than pure B_z.
B_{xz}	Mixed case with out-of-plane component. Gives one of the largest final interface heights and a clearly broadened head.
B_{xyz}	Fully mixed direction. Intermediate behavior between the in-plane and strongly out-of-plane cases.

kinetic channel, so the solution stays closer to the kinetic reference. When $\mathbf{B}$ is rotated toward z, less of the in-plane RTI dynamics is carried by that kinetic channel and more of it is handled through the reduced perpendicular response. The result is a slightly faster instability and a broader, more mushroom-like interface. From that point of view, the out-of-plane cases behave more fluid-like in their in-plane evolution, whereas the in-plane cases retain more of the kinetic moderation seen in the fully kinetic runs.

3.3.2 Directional transport terms and their relation to PKPM field orientation

The directional transport data clarify what the imposed field orientation is doing in PKPM. Since these are neutral runs, the field does not enter through a physical Lorentz force. Instead, the imposed $\hat{\mathbf{b}}$ selects the field-aligned kinetic direction, so changing the field orientation changes which part of the in-plane RTI dynamics is treated through kinetic transport and which part is left to the perpendicular moment closure.

Table 3.6: Frame-10 directional heat-flux magnitudes and ideal-to-nonideal ratios for the fully kinetic reference and PKPM orientation cases, together with the final interface height and fitted growth rate from the corresponding evolution analysis. Here $R_x = \langle|\text{ideal}_x|\rangle/\langle|\text{nonideal}_x|\rangle$ and $R_y = \langle|\text{ideal}_y|\rangle/\langle|\text{nonideal}_y|\rangle$.

Case	$\langle\lvert q_x\rvert\rangle$	$\langle\lvert q_y\rvert\rangle$	R_x	R_y	$h_x(t/\tau_{\mathrm{RT}}=3)$	γ
Kinetic LBO	8.95×10^{-4}	6.81×10^{-4}	76.17	45.92	0.9937	0.7779
PKPM B_x	1.83×10^{-3}	0	47.17	4974.16	1.0978	0.7893
PKPM B_y	0	1.43×10^{-3}	4384.98	25.06	1.0691	0.7936
PKPM B_z	0	0	93520.97	73750.13	1.1655	0.8038
PKPM B_{xy}	1.29×10^{-3}	1.29×10^{-3}	66.80	28.68	1.1064	0.7887
PKPM B_{yz}	0	7.55×10^{-4}	7627.52	52.20	1.1329	0.8014
PKPM B_{xz}	9.63×10^{-4}	0	98.33	8232.35	1.1376	0.7922
PKPM B_{xyz}	8.77×10^{-4}	8.77×10^{-4}	105.01	43.99	1.1308	0.7935

The first important result is that the non-ideal contribution is controlled mainly by the heat-flux term, not by the stress-work term. In all cases where an in-plane heat-flux component is present, $|q_i|$ is much larger than $|u_j\Pi_{ji}|$. The stress-work term is therefore a

correction, whereas the heat flux sets the main non-ideal transport channel. This means that the field-orientation dependence is telling us, above all, how PKPM is redistributing thermal energy within the x–y plane.

The second important result is that the presence or absence of q_x and q_y follows the field direction very clearly. Cases with an x-component in $\hat{\mathbf{b}}$ (B_x, B_{xy}, B_{xz}, B_{xyz}) produce nonzero q_x. Cases with a y-component (B_y, B_{xy}, B_{yz}, B_{xyz}) produce nonzero q_y. The pure B_z case produces essentially zero in-plane heat flux in both directions. This is exactly the behavior expected from the PKPM construction: the dominant heat transport is carried along the field-aligned kinetic direction, so only the in-plane projection of that direction appears in the resolved q_x and q_y components.

Before comparing the PKPM q_y structure across different field orientations, it is useful to compare the signed q_y field directly between the fully kinetic reference and the PKPM B_{xy} case, since B_{xy} is the simplest PKPM orientation that retains both in-plane transport channels.

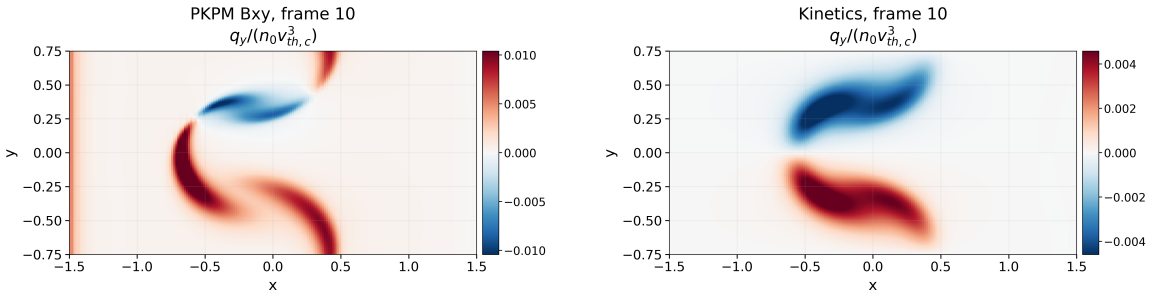

(a) PKPM B_{xy}, signed q_y at frame 10.

(b) Fully kinetic reference, signed q_y at frame 10.

Figure 3.9: Comparison of q_y at frame 10 for the PKPM B_{xy} case and the fully kinetic reference. Both cases show oppositely signed transport along the upper and lower shoulders of the interface, while the kinetic field remains broader and smoother than the PKPM projection.

A comparison of the signed q_y field between the fully kinetic reference and the PKPM B_{xy} case helps clarify both what PKPM captures and what it misses. In the kinetic solution, q_y forms a broad bipolar structure along the upper and lower shoulders of the interface,

showing that thermal energy is redistributed in opposite y-directions above and below the centerline. The PKPM B_{xy} case reproduces the same qualitative pattern, with oppositely signed q_y lobes appearing along the two shoulders, so the reduced model does retain the main shoulder-directed transport trend. At the same time, the PKPM pattern is sharper and more localized than in the kinetic case. This difference is consistent with the lowest-order PKPM closure: in the B_{xy} case, the plotted q_x and q_y are not independent heat-flux components, but equal in-plane projections of the same retained field-aligned reduced heat flux. The comparison is therefore useful because it shows that PKPM captures the dominant signed transport structure qualitatively, while the fully kinetic model retains a broader transverse heat-flux structure that is not explicitly represented in PKPM at this order.

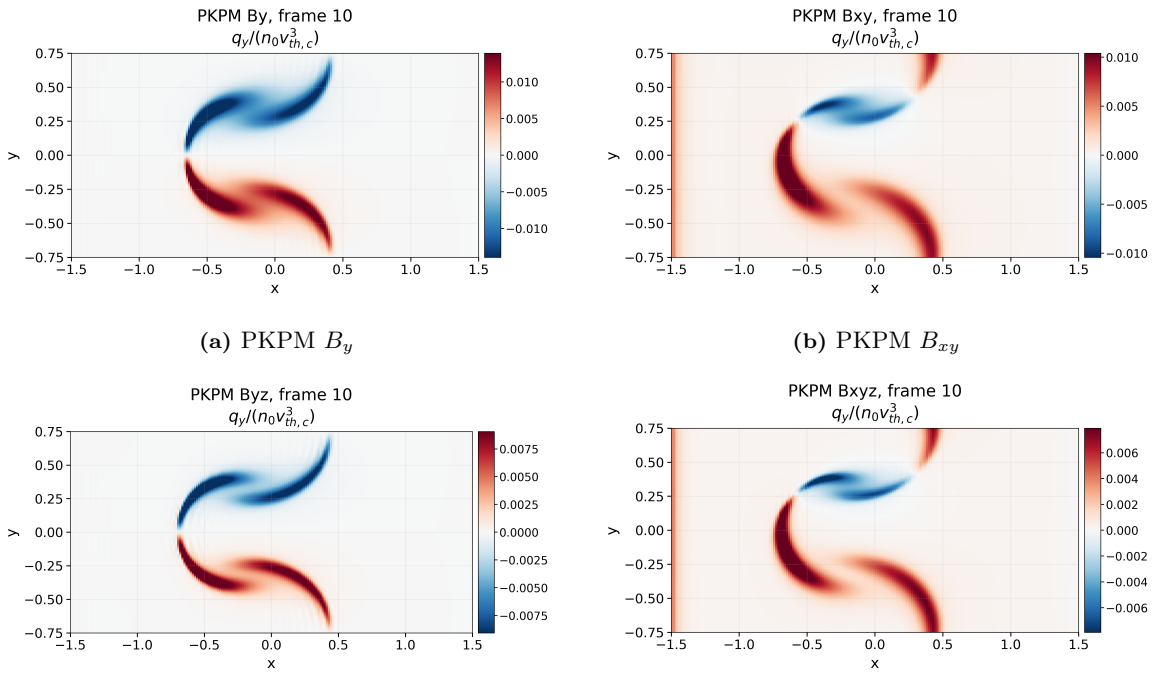

(a) PKPM B_y

(b) PKPM B_{xy}

(c) PKPM B_{yz}

(d) PKPM B_{xyz}

Figure 3.10: Signed q_y at frame 10 for PKPM cases with different field orientations. Cases with a nonzero b_y component show clear red-blue structures along the upper and lower interface shoulders, indicating oppositely directed transverse heat-flux transport above and below the centerline. In contrast, the B_z case shows essentially no organized q_y transport in the resolved plane.

When $b_y \neq 0$, PKPM permits a nonzero q_y through the field-aligned projection

$$q_y \approx (q_\parallel + q_\perp)\, b_y.$$

Thus, the red-blue structures in the signed q_y plots are showing oppositely directed y-transport. Because the upper and lower RTI shoulders are tilted in opposite y-directions, a field-aligned transport channel with a y-projection naturally appears with opposite sign above and below the centerline. This behavior is shown directly in Fig. 3.10. In this sense, the y-component of $\hat{\mathbf{b}}$ redirects part of the non-ideal energy transport along the interface shoulders, rather than primarily into the head-advance direction x. This is a direct consequence of the PKPM closure: changing $\hat{\mathbf{b}}$ changes the direction along which the retained kinetic heat-flux response is projected into the resolved plane.

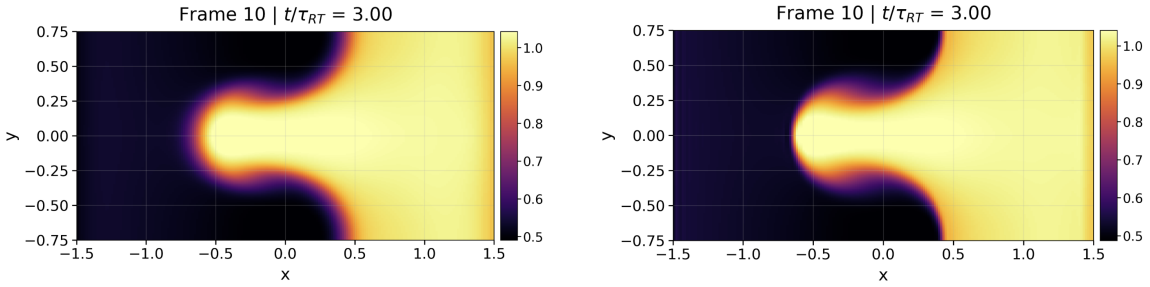

(a) PKPM with $\mathbf{B} \parallel \hat{\mathbf{x}}$, frame 10.

(b) PKPM with $\mathbf{B} \parallel \hat{\mathbf{y}}$, frame 10.

Figure 3.11: Final density field at $t/\tau_{\mathrm{RT}} = 3.00$ for the PKPM B_x and B_y cases. The overall RTI morphology is similar in both runs, but the interface head and the upper and lower shoulders show small geometric differences.

This also helps explain why the ordering by fitted growth rate is not identical to the ordering by final interface height. The clearest example is B_y. In that case, the fitted growth rate is slightly larger than for B_x and B_{xy}, but the final interface height is smaller. The reason is that the fitted γ measures the growth of h_x only over the selected quasi-exponential interval, whereas the final h_x also includes later nonlinear reshaping of the interface. For B_y, the dominant non-ideal transport is in the y-direction rather than the x-direction:

$$\langle |q_x| \rangle \approx 0, \qquad \langle |q_y| \rangle = 1.42855 \times 10^{-3}.$$

Figure 3.11 compares the final density fields for the PKPM B_x and B_y cases. It shows that the B_x case has a more diffuse interface than the B_y case. This is consistent with the transport data. In B_x, the dominant non-ideal transport is carried by q_x, which acts directly across the interface and therefore smooths the front more strongly. In B_y, the main non-ideal transport is instead carried by q_y, which redistributes energy along the upper and lower shoulders rather than across the interface itself. As a result, B_x shows stronger cross-interface diffusion, while B_y mainly modifies the flanks. This also helps explain why B_{xy} remains the closest PKPM case to the fully kinetic reference from a transport perspective, since it retains both in-plane transport channels. This also means that once the interface shoulders develop, PKPM is preferentially redistributing energy along the upper and lower flanks instead of feeding continued extension of the head in the x-direction. As a result, the mode can maintain a competitive linear-stage growth rate while finishing with a smaller overall x-extent.

The opposite limit is B_z. In that case both q_x and q_y are essentially absent, so the in-plane dynamics become the most ideal-like of all the PKPM runs. This is exactly the case with the largest fitted growth rate and the largest final interface height, and it is also the case with the clearest mushroom-like head. The interpretation is that once the in-plane non-ideal redistribution is strongly reduced, the interface is freer to roll up and advance without being smoothed by in-plane heat transport.

The mixed z-containing cases B_{yz} and B_{xz} sit between these two limits. Including a z-component weakens the in-plane heat flux relative to the purely in-plane cases, so both cases grow more strongly than B_x, B_y, and B_{xy}. At the same time, they retain one in-plane transport channel, so they do not become as extreme as the pure B_z case. This is why they cluster just below B_z in growth and late-time interface extent.

The fully kinetic case is also helpful here. Unlike PKPM, the fully kinetic run retains both in-plane transport channels at once, with moderate nonzero values of both q_x and q_y. Among the PKPM cases, B_{xy} is therefore the closest analogue in transport structure, because it is the only simple orientation that retains both in-plane channels at comparable strength. That is why B_{xy} remains one of the closest PKPM cases to the kinetic reference in overall morphology, even though its fitted growth rate is not the largest.

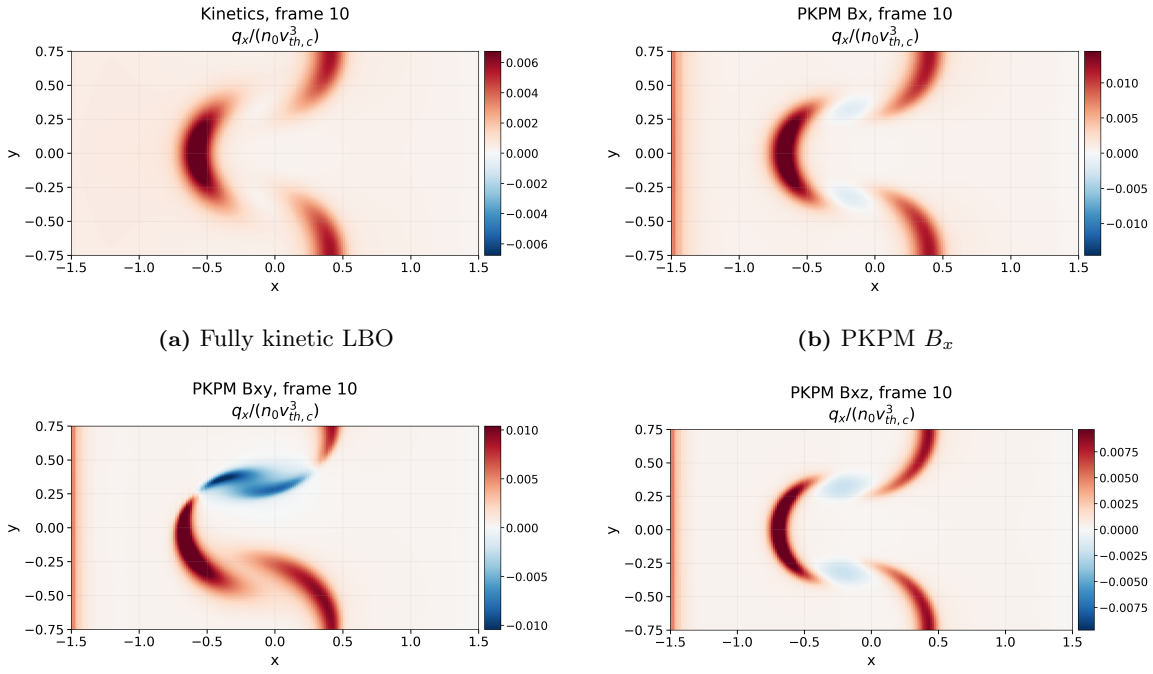

(a) Fully kinetic LBO

(b) PKPM B_x

(c) PKPM B_{xy}

(d) PKPM B_{xz}

Figure 3.12: q_x at frame 10 for the fully kinetic LBO case and selected PKPM field orientations. These panels show how the head-advance heat-flux channel changes when the imposed field retains or loses an x-component. Cases with nonzero b_x preserve a strong forward transport pathway in the growth direction, while differences in the detailed lobe structure reflect how the PKPM kinetic direction redistributes non-ideal transport across the interface head and shoulders.

Table 3.7: Frame-10 heat-flux and stress-work contributions to the non-ideal transport in the x- and y-directions. The data show that whenever an in-plane heat-flux channel is present, it dominates the non-ideal contribution, while the stress-work term remains much smaller.

Case	$\lvert q_x \rvert$	$\lvert u_j \Pi_{jx} \rvert$	$\lvert q_y \rvert$	$\lvert u_j \Pi_{jy} \rvert$
Kinetic LBO	8.95×10^{-4}	4.24×10^{-5}	6.81×10^{-4}	1.21×10^{-5}
PKPM B_x	1.83×10^{-3}	3.87×10^{-5}	0	7.11×10^{-6}
PKPM B_y	0	2.02×10^{-5}	1.43×10^{-3}	1.57×10^{-5}
PKPM B_z	0	1.07×10^{-6}	0	5.72×10^{-7}
PKPM B_{xy}	1.29×10^{-3}	1.73×10^{-5}	1.29×10^{-3}	2.05×10^{-5}
PKPM B_{yz}	0	1.26×10^{-5}	7.55×10^{-4}	2.24×10^{-6}
PKPM B_{xz}	9.63×10^{-4}	5.74×10^{-6}	0	4.71×10^{-6}
PKPM B_{xyz}	8.77×10^{-4}	7.93×10^{-6}	8.77×10^{-4}	9.75×10^{-6}

Taken together, the directional transport data show that the orientation dependence in PKPM is not arbitrary. Cases with strong in-plane heat-flux transport are more strongly redistributed within the simulation plane, whereas cases with weak in-plane heat flux behave more ideally and develop stronger late-time RTI roll-up. In this sense, the field direction in PKPM chooses direction in which kinetic transport is allowed to smooth or redistribute the instability.

3.4 Dimensionality sensitivity of the kinetic reference

A natural question in the present comparison is whether using a $2x2v$ kinetic reference could affect the interpretation of the PKPM results. To check this, an auxiliary BGK comparison was carried out between $2x2v$ and $2x3v$ kinetic runs. This is not meant to replace the main $2x2v$ LBO reference, but rather to estimate how sensitive the kinetic solution is to the addition of a third velocity dimension.

A full $2x3v$ LBO reference was not practical at the resolution used here because the computational cost grows very rapidly with phase-space dimensionality. In the present DG discretization, the $2x2v$ case uses 48 basis functions, whereas the $2x3v$ case uses 112. Since the cost scales approximately with the square of the number of basis functions, and the

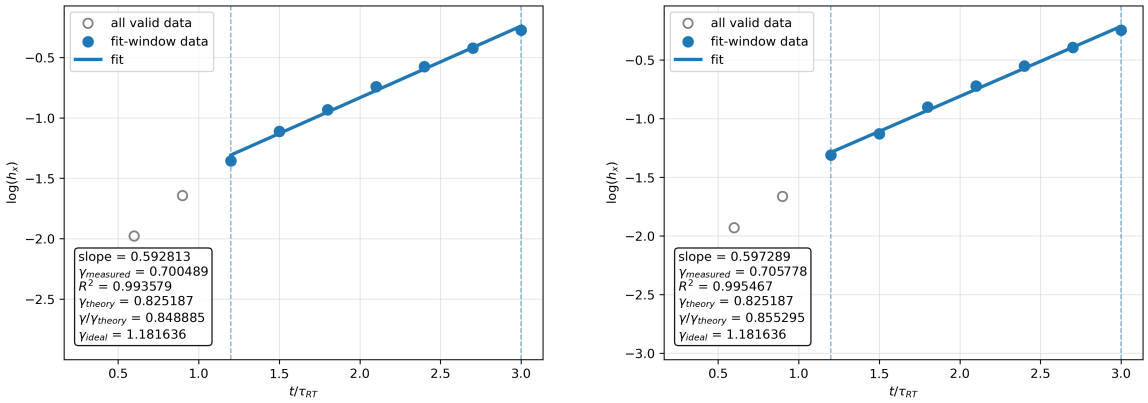

(a) BGK kinetic $2x2v$.

(b) BGK kinetic $2x3v$.

Figure 3.13: Growth-rate fits for the BGK kinetic runs used to assess dimensionality sensitivity. Adding the third velocity dimension changes the fitted growth rate only slightly, indicating that the $2x2v$ reference captures the dominant growth behavior relevant to the PKPM comparison.

extra velocity dimension also increases the phase-space workload, the overall cost increase is roughly 80 times the 2x2v LBO case. In practice, this means that a comparable $2x3v$ LBO run would be expensive.

As shown in Fig. 3.13, the fitted growth rates for the $2x2v$ and $2x3v$ BGK runs are very close, with only a small increase in γ when the third velocity dimension is included. The final interface height increases from $h_x(t/\tau_{\mathrm{RT}} = 3) = 0.7599$ in $2x2v$ to 0.7821 in $2x3v$, which is a difference of about 2.9%. The fitted growth rate changes from $\gamma = 0.7005$ to $\gamma = 0.7058$, corresponding to an increase of only about 0.8%. The density evolution remains qualitatively the same in both cases, with the $2x3v$ run showing only a slightly larger late-time interface extension.

This indicates that the extra velocity dimension introduces a modest correction, but does not change the overall RTI behavior. If we expand what was observed in BGK case, the $2x2v$ kinetic reference appears to capture the dominant dynamics relevant to the PKPM comparison, while a fully resolved $2x3v$ LBO reference would likely refine the result slightly only by 1% at much greater computational cost.

3.5 Magnetic-field orientation effects at Kn = 0.1

Figure 3.14 compares the final density fields at $t/\tau_{\mathrm{RT}} = 3$ for the fully kinetic reference and all PKPM field-direction cases. Compared with the Kn = 0.01 results, the Kn = 0.1 regime is much more transport-dominated. The fully kinetic solution already shows substantial interface diffusion, a weak head, and no strong mushroom roll-up. This indicates that at intermediate collisionality, non-ideal transport is no longer a small correction; it directly competes with the RTI growth.

For the neutral PKPM model, the imposed $\mathbf{B}$ defines the PKPM parallel direction $\hat{\mathbf{b}}$, which determines how the field-aligned kinetic transport projects into the 2D (x, y) plane. In simple terms, the in-plane transport channels scale with the in-plane projection of $\hat{\mathbf{b}}$,

$$q_x \sim b_x q_{\parallel}, \qquad q_y \sim b_y q_{\parallel},$$

so adding a z-component reduces the part of the parallel transport that can act inside the simulation plane. Thus, when b_z increases, less of the field-aligned non-ideal transport is available to smooth or redistribute the interface in x and y. It means that the PKPM closure contributes less in-plane redistribution relative to the ideal RTI advection/compression terms. As a result, the density contour is less smeared by transport and is freer to preserve a sharper RTI structure.

Table 3.8: Frame-10 transport budget in the x-direction for the fully kinetic reference and PKPM Kn = 0.1 cases.

Case	$\langle\lvert I_1\rvert\rangle$	$\langle\lvert I_2\rvert\rangle$	$\langle\lvert q_x\rvert\rangle$	$\langle\lvert\text{stress}\rvert\rangle$	$\langle\lvert\text{ideal}\rvert\rangle$	$\langle\lvert\text{nonideal}\rvert\rangle$	ratio
Kinetic	2.13287×10^{-2}	2.47357×10^{-6}	9.02130×10^{-3}	4.34989×10^{-6}	2.13312×10^{-2}	9.02565×10^{-3}	2.36340
PKPM B_x	5.39752×10^{-2}	7.41011×10^{-5}	1.78958×10^{-2}	6.39439×10^{-5}	5.40493×10^{-2}	1.79597×10^{-2}	3.00947
PKPM B_y	4.35395×10^{-2}	4.21556×10^{-5}	0	2.90339×10^{-5}	4.35816×10^{-2}	2.90339×10^{-5}	1501.06305
PKPM B_z	1.01030×10^{-1}	9.95483×10^{-4}	0	9.54872×10^{-6}	1.02026×10^{-1}	9.54872×10^{-6}	10684.77930
PKPM B_{xy}	4.20964×10^{-2}	6.86843×10^{-5}	9.20231×10^{-3}	4.22542×10^{-5}	4.21651×10^{-2}	9.24456×10^{-3}	4.56107
PKPM B_{yz}	7.50260×10^{-2}	3.56804×10^{-4}	0	3.92645×10^{-5}	7.53828×10^{-2}	3.92645×10^{-5}	1919.87018
PKPM B_{xz}	7.49811×10^{-2}	3.21014×10^{-4}	8.94117×10^{-3}	2.12740×10^{-5}	7.53021×10^{-2}	8.96244×10^{-3}	8.40196
PKPM B_{xyz}	5.95452×10^{-2}	2.36911×10^{-4}	6.94920×10^{-3}	2.99699×10^{-5}	5.97821×10^{-2}	6.97917×10^{-3}	8.56580

The transport budgets in Tables 3.8 and 3.9 make this trend clear. In the fully kinetic

case, the non-ideal terms are already comparable to the ideal contribution:

$$\frac{\langle|\text{ideal}_x|\rangle}{\langle|\text{nonideal}_x|\rangle} = 2.36, \qquad \frac{\langle|\text{ideal}_y|\rangle}{\langle|\text{nonideal}_y|\rangle} = 3.72.$$

So the kinetic $\text{Kn} = 0.1$ reference is strongly diffusive.

Among the PKPM cases, the strongest roll-up appears for B_z. This is also the clearest case where the in-plane non-ideal transport is nearly removed:

$$\langle|\text{nonideal}_x|\rangle = 9.55 \times 10^{-6}, \qquad \langle|\text{nonideal}_y|\rangle = 2.93 \times 10^{-6},$$

with correspondingly huge ideal-to-nonideal ratios in both directions. So B_z is not "more unstable" because of magnetic forcing; rather, it is the case where the in-plane closure-driven smoothing is weakest. Once that smoothing is minimized, the interface can maintain a much sharper head and develop the strongest mushroom-like morphology.

The B_x and B_{xz} pair shows the same mechanism in a more controlled way. In both cases, the active non-ideal transport is mainly in the x-direction, but adding the z-component reduces its in-plane effect. For B_x,

$$\langle|\text{nonideal}_x|\rangle = 1.796 \times 10^{-2}, \qquad \frac{\langle|\text{ideal}_x|\rangle}{\langle|\text{nonideal}_x|\rangle} = 3.01,$$

whereas for B_{xz},

$$\langle|\text{nonideal}_x|\rangle = 8.962 \times 10^{-3}, \qquad \frac{\langle|\text{ideal}_x|\rangle}{\langle|\text{nonideal}_x|\rangle} = 8.40.$$

So in this family, adding z weakens the x-directed smoothing both in absolute terms and relative to the ideal transport. The density field then becomes more RTI-like, with a clearer forward head.

The B_y and B_{yz} cases highlight an important nuance. Here the dominant non-ideal channel is in the y-direction, so the transport acts mainly along the upper and lower shoulders rather than directly in the head-advance direction. For B_y,

$$\langle|\text{nonideal}_y|\rangle = 5.701 \times 10^{-3}, \qquad \frac{\langle|\text{ideal}_y|\rangle}{\langle|\text{nonideal}_y|\rangle} = 2.67.$$

For B_{yz},

$$\langle|\text{nonideal}_y|\rangle = 5.892 \times 10^{-3}, \qquad \frac{\langle|\text{ideal}_y|\rangle}{\langle|\text{nonideal}_y|\rangle} = 5.39.$$

Table 3.9: Frame-10 transport budget in the y-direction for the fully kinetic reference and PKPM Kn = 0.1 cases.

Case	$\langle\lvert I_1\rvert\rangle$	$\langle\lvert I_2\rvert\rangle$	$\langle\lvert q_y\rvert\rangle$	$\langle\lvert\text{stress}\rvert\rangle$	$\langle\lvert\text{ideal}\rvert\rangle$	$\langle\lvert\text{nonideal}\rvert\rangle$	ratio
Kinetic	2.93035×10^{-3}	2.40387×10^{-7}	7.81387×10^{-4}	5.93053×10^{-6}	2.93059×10^{-3}	7.87318×10^{-4}	3.72225
PKPM B_x	1.45860×10^{-2}	7.73583×10^{-6}	0	6.79462×10^{-6}	1.45937×10^{-2}	6.79462×10^{-6}	2147.83488
PKPM B_y	1.52196×10^{-2}	6.82664×10^{-6}	5.68768×10^{-3}	1.34671×10^{-5}	1.52265×10^{-2}	5.70114×10^{-3}	2.67077
PKPM B_z	4.32193×10^{-2}	2.06483×10^{-4}	0	2.92934×10^{-6}	4.34258×10^{-2}	2.92934×10^{-6}	14824.43238
PKPM B_{xy}	2.58892×10^{-2}	2.60738×10^{-5}	9.20231×10^{-3}	4.28993×10^{-5}	2.59153×10^{-2}	9.24520×10^{-3}	2.80311
PKPM B_{yz}	3.16662×10^{-2}	7.33885×10^{-5}	5.88860×10^{-3}	3.88240×10^{-6}	3.17396×10^{-2}	5.89248×10^{-3}	5.38645
PKPM B_{xz}	2.89594×10^{-2}	5.30616×10^{-5}	0	1.30710×10^{-5}	2.90125×10^{-2}	1.30710×10^{-5}	2219.61110
PKPM B_{xyz}	3.04585×10^{-2}	6.25258×10^{-5}	6.94920×10^{-3}	2.72785×10^{-5}	3.05211×10^{-2}	6.97648×10^{-3}	4.37485

So the absolute y-flux is not reduced much by adding z, but its *dynamical importance* is reduced because the ideal transport grows more strongly. This is why adding a z-component should be interpreted as reducing the *effect* of the non-ideal transport inside the plane, not necessarily lowering every non-ideal number by itself.

This also helps explain why morphology is not controlled by fitted growth rate alone. A case such as B_y can still maintain a competitive growth during the quasi-exponential stage, but because its dominant transport acts along y, it redistributes energy along the two tilted shoulders instead of feeding continued forward extension of the head. The result is a smaller final x-extent, even though the fitted γ is not especially low. In other words, the linear-stage growth rate measures how fast the interface amplitude grows over the selected interval, while the final frame also includes later nonlinear reshaping by directional transport.

The mixed cases B_{xy} and B_{xyz} show both effects at once. For B_{xy}, the non-ideal transport is strong in both directions,

$$\langle\lvert\text{nonideal}_x\rvert\rangle = 9.245 \times 10^{-3}, \qquad \langle\lvert\text{nonideal}_y\rvert\rangle = 9.245 \times 10^{-3},$$

so the interface is smoothed both in the head-advance direction and along the shoulders. When a z-component is added, giving B_{xyz},

$$\langle\lvert\text{nonideal}_x\rvert\rangle = 6.979 \times 10^{-3}, \qquad \langle\lvert\text{nonideal}_y\rvert\rangle = 6.976 \times 10^{-3},$$

and the ideal-to-nonideal ratios increase in both directions. Thus, adding z weakens the

Table 3.10: Effect of adding a z-component to the PKPM field direction at Kn = 0.1. The comparison is made within each family.

Family	Change in dominant non-ideal channel	Change in ratio	Interpretation
$B_x \to B_{xz}$	$q_x : 1.796 \times 10^{-2} \to 8.962 \times 10^{-3}$	$R_x : 3.01 \to 8.40$	x-directed smoothing weakens strongly.
$B_y \to B_{yz}$	$q_y : 5.701 \times 10^{-3} \to 5.892 \times 10^{-3}$	$R_y : 2.67 \to 5.39$	Absolute q_y stays similar, but its effect weakens relative to the ideal transport.
$B_{xy} \to B_{xyz}$	$q_x, q_y : 9.245 \times 10^{-3} \to 6.98 \times 10^{-3}$	$R_x : 4.56 \to 8.57,\ R_y : 2.80 \to 4.37$	Both in-plane smoothing channels weaken.

in-plane transport channels in both absolute and relative terms. The final density contour therefore becomes more RTI-like than B_{xy}, but still not as rolled up as the pure B_z limit, because some in-plane smoothing remains active.

Taken together, the Kn = 0.1 results show that the field orientation in PKPM acts as a geometric control on *where* non-ideal transport is allowed to act in the 2D plane. Cases with an x-component preferentially smooth the advancing head, cases with a y-component redistribute energy along the upper and lower shoulders, cases with both components combine both effects, and cases with a z-component reduce the in-plane influence of the field-aligned transport and therefore recover a sharper, more RTI-like interface. In x-dominated families this happens partly because the in-plane non-ideal term decreases in absolute magnitude, whereas in y-dominated families it happens mainly because the ideal term grows while the non-ideal term stays similar.

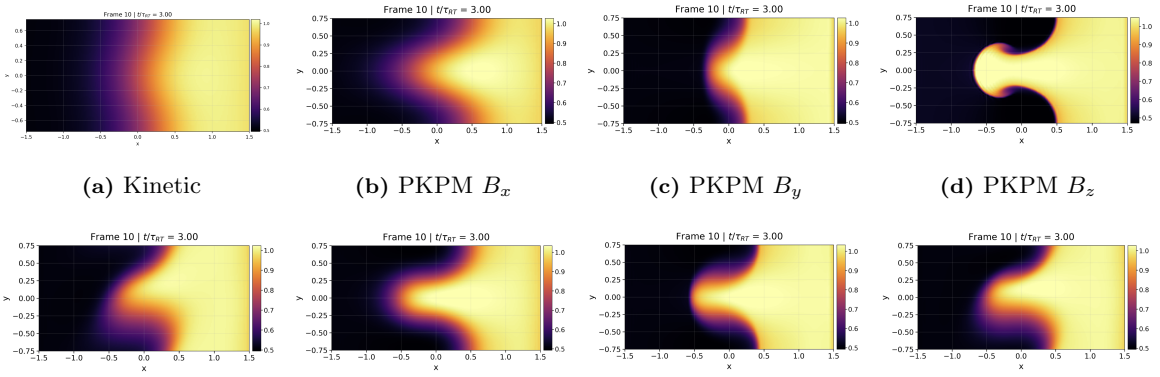

(a) Kinetic **(b)** PKPM B_x **(c)** PKPM B_y **(d)** PKPM B_z

(e) PKPM B_{xy} **(f)** PKPM B_{xz} **(g)** PKPM B_{yz} **(h)** PKPM B_{xyz}

Figure 3.14: Final density evolution at $t/\tau_{\mathrm{RT}} = 3$ for the fully kinetic $\mathrm{Kn} = 0.1$ reference and the PKPM cases with different imposed field directions.

3.6 Computational cost and fidelity

In this work, the fully kinetic runs are treated as the reference, and PKPM is assessed by asking two practical questions: how much cheaper is it, and how closely does it reproduce the kinetic RTI behavior that matters for this study?

The computational difference is large. For the matched $\mathrm{Kn} = 0.01$ LBO case, the fully kinetic $2x2v$ run required about 25 hours on 192 cores, or roughly 4800 core-hours. The corresponding PKPM LBO runs required only about 35 to 45 minutes on the same 192 cores, or about 112 to 144 core-hours. This is a cost reduction of roughly a factor of 33 to 43. The higher cost of the kinetic model comes from resolving the full phase-space dynamics, and it increases further when either the collision physics or the dimensionality is made more complete. This is also clear from the auxiliary kinetic runs: BGK $2x2v$ required only about 2 hours on 128 cores, whereas BGK $2x3v$ required about 1 day 14 hours on the same core count, and the LBO $2x2v$ reference cost about the same total core-hours as BGK $2x3v$. In other words, both the collision operator and the velocity-space dimensionality matter strongly for cost, and either one can make the kinetic reference expensive.

What makes PKPM useful is that the loss in fidelity is much smaller than the savings in cost. For the $\mathrm{Kn} = 0.01$ LBO comparison, the kinetic reference gives $\gamma = 0.7779$ and

Table 3.11: Representative computational cost of the kinetic and PKPM runs used for comparison.

Case	Wall time	Core-hours
Kinetic LBO Kn = 0.01, $2x2v$	25 h	4800
PKPM LBO Kn = 0.01	35–45 min	112–144
Kinetic LBO Kn = 0.1	1 h 57 min	499
PKPM LBO Kn = 0.1	10–16 min	11–17

$h_x(t/\tau_{\mathrm{RT}} = 3) = 0.9937$. Across the PKPM field-direction cases, the fitted growth rates lie between 0.7887 and 0.8038, so the difference from the kinetic value is only about 1.4% to 3.3%. The final interface heights differ more, ranging from 7.6% to 17.3% above the kinetic result, which shows that the later nonlinear morphology is more sensitive to the PKPM closure direction than the fitted linear-stage growth. Even so, the main kinetic trends are still captured at a much lower computational cost.

Overall, the fully kinetic model remains the higher-fidelity reference, but PKPM reproduces the main RTI trends well enough to be useful at much lower computational cost. The agreement is stronger at Kn = 0.01, where PKPM remains relatively close to the kinetic growth and morphology, whereas at Kn = 0.1 the deviation from the strongly diffusive kinetic solution become more pronounced. This suggests that, for the present problem, PKPM captures the kinetic dynamics more faithfully in the intermediate collisional Kn = 0.01 regime than in the low collisional Kn = 0.1 regime.

Chapter 4

CONCLUSION AND FUTURE WORK

This thesis compares the Rayleigh-Taylor instability in a fully kinetic model and in the parallel-kinetic-perpendicular-moment (PKPM) model in order to understand what the reduced description captures, where it begins to differ from the kinetic reference, and how those differences depend on collisionality and on the direction of the retained kinetic response. The comparison is built on matched initial conditions, equivalent perturbations, and a common interface-tracking procedure, so that the later differences between the models can be interpreted as true modeling effects rather than setup inconsistencies.

For the matched $\mathrm{Kn} = 0.01$ LBO case, PKPM reproduces the main features of the fully kinetic reference reasonably well. The density evolution, interface displacement, and fitted growth rate remain close to the kinetic solution, showing that the reduced model can capture the dominant instability behavior. At the same time, the later nonlinear interface shape depends more strongly on the PKPM field direction than the fitted linear-stage growth. The in-plane cases remain closest to the kinetic reference, while cases containing a z-component develop broader heads and more mushroom-like roll-up. This shows that the field direction in PKPM is not just a numerical parameter. It changes how the model partitions the dynamics between the retained kinetic direction and the reduced perpendicular response, and that choice leaves a clear signature in the morphology of the instability.

The transport analysis helps explain why these differences appear. In the $\mathrm{Kn} = 0.01$ cases, the ideal flux remains dominant, but the non-ideal contribution is controlled mainly by heat flux rather than stress work. The directional heat-flux channels follow the PKPM field direction in a clear way: cases with an x-component retain non-ideal transport in the head-advance direction, cases with a y-component redistribute energy along the upper and lower shoulders, and the pure B_z case removes most of the in-plane heat-flux response.

At $\mathrm{Kn} = 0.1$, these directional effects become even more visible. The fully kinetic

reference is strongly diffusive and does not develop a clear RTI head over the simulated interval. The PKPM cases, however, show a much wider spread of late-time behavior. Some orientations remain more diffusive, some redistribute energy mainly along the shoulders, and some recover a much sharper and more RTI-like head. In particular, adding a z-component reduces the influence of the in-plane non-ideal transport and allows a clearer roll-up to form. This makes the Kn $= 0.1$ study especially useful, because it shows that the PKPM field direction can change whether the interface remains transport-dominated or develops a recognizable RTI structure at all. The BGK $2x2v$ and $2x3v$ comparison is used to check whether a missing third velocity dimension would significantly affect the reference, and the small difference between them supports using $2x2v$ LBO as a reasonable kinetic benchmark.

A final outcome of this work is the cost–fidelity balance between the two model classes. The fully kinetic model remains the more complete description, but it is substantially more expensive. PKPM, by contrast, captures the main growth behavior and the main transport-driven morphology trends at a much lower computational cost. The conclusion is therefore not that PKPM replaces the kinetic model, but that it provides a useful reduced description when the goal is to capture the dominant instability physics without resolving the full phase space. Which model is more appropriate depends on what one wants from the simulation: if the full velocity-space response and non-equilibrium transport must be resolved directly, the kinetic model is the more faithful choice; if the main interest is in instability growth, field-direction effects, and transport-driven changes in morphology at a manageable cost, PKPM becomes a practical and informative alternative.

Future work should extend the present comparison beyond the neutral-species setting used in this thesis and into genuinely magnetized plasma RTI, where the magnetic field would influence not only the geometric partition of the PKPM closure but also the instability physics itself through self-consistent electromagnetic evolution. A natural next step is therefore to study charged multispecies plasmas, so that the comparison can move closer to the weakly collisional magnetized environments for which PKPM is designed, and so that the roles of field strength, field shear, anisotropic transport, and finite collisionality in shaping RTI growth and nonlinear morphology can be examined more directly. It would also be important to move beyond the spatially uniform collision setting used here and allow the col-

lisionality to vary in space, since in more realistic plasmas the collision frequency depends on local density, temperature, composition, and ionization state. That extension would make it possible to study how collisional smoothing and kinetic effects vary across the interface as the instability evolves, rather than treating the entire domain with a single collisional background. More broadly, the framework developed here is relevant beyond the specific RTI problem studied in this thesis. Weakly collisional magnetized plasmas often exhibit strongly anisotropic transport, and reduced models such as PKPM are especially attractive in such settings because they retain field-aligned kinetic response without requiring the full cost of a continuum kinetic treatment in all directions. This makes the present comparison framework potentially useful for a broader class of simulations in which perpendicular transport, finite Larmor radius effects, or field-aligned kinetic dynamics play an important role. Examples include low-β magnetized plasmas in confinement devices such as tokamaks and stellarators, as well as sheared-flow-stabilized Z-pinches, where transport across the field and the competition between ideal and non-ideal effects are central to the overall plasma evolution [5]. A further development would be to include multispecies collisional physics and eventually partially ionized or two-fluid effects, which would allow ion-electron coupling, ion-neutral decoupling, ionization and recombination, and frictional heating to be captured within the same framework. In this way, the next stage of the work would not only refine the benchmark developed here, but would move the kinetic-PKPM comparison toward a more complete plasma RTI problem in which both the reduced model and the kinetic reference are tested against genuinely magnetized, multispecies, and spatially nonuniform dynamics.

BIBLIOGRAPHY

[1] Colin S. Adams, Auna L. Moser, and Scott C. Hsu. Observation of Rayleigh–Taylor-instability evolution in a plasma with magnetic and viscous effects. *Physical Review E*, 92(5):051101, 2015.

[2] Ratan Kumar Bera, Yang Song, and Bhuvana Srinivasan. The effect of viscosity and resistivity on rayleigh–taylor instability induced mixing in magnetized high-energy-density plasmas. *Journal of Plasma Physics*, 88(2):905880209, 2022.

[3] 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. *Physical Review*, 94(3):511–525, 1954.

[4] Petr Cagas, Ammar H. Hakim, and Bhuvana Srinivasan. A boundary value "reservoir problem" and boundary conditions for multi-moment multifluid simulations of sheaths. *Physics of Plasmas*, 28(1), 2021.

[5] Jack Coughlin, Jingwei Hu, and Uri Shumlak. Asymptotic perpendicular transport in low-beta collisionless plasma, 2024.

[6] J. P. Dougherty. Model Fokker–Planck equation for a plasma and its solution. *Physics of Fluids*, 7(11):1788–1799, 1964.

[7] R. E. Duff, F. H. Harlow, and C. W. Hirt. Effects of diffusion on interface instability between gases. *Physics of Fluids*, 5(4):417–425, 1962.

[8] Gkeyll Developers. Gkeyll 1.0 documentation. `https://gkeyll.readthedocs.io/`. Accessed: 2026-05-10.

[9] Gkeyll Developers. Publications and theses — Gkeyll 1.0 documentation. `https://gkeyll.readthedocs.io/en/latest/gkeyll/pubs.html`. Accessed: 2026-05-10.

[10] Robert J. Goldston and Paul H. Rutherford. *Introduction to Plasma Physics*. Institute of Physics Publishing, Bristol and Philadelphia, 1995.

[11] Ryosuke Goto, Hideaki Miura, Atsushi Ito, Masahiko Sato, and Tomoharu Hatori. Hall and gyro-viscosity effects on the Rayleigh–Taylor instability in a 2d rectangular slab. *Plasma and Fusion Research*, 9:1403076, 2014.

[12] Ammar Hakim, Manaure Francisquez, James Juno, and Gregory W. Hammett. Conservative discontinuous Galerkin schemes for nonlinear Dougherty–Fokker–Planck collision operators. *Journal of Plasma Physics*, 86(4):905860403, 2020.

[13] 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. *Physical Review Letters*, 59(26):2971–2974, 1987.

[14] Atsushi Ito and Hideaki Miura. Parameter dependence of two-fluid and finite Larmor radius effects on the Rayleigh–Taylor instability in finite beta plasmas. *Physics of Plasmas*, 23(12):122123, 2016.

[15] J. Juno, A. Hakim, J. TenBarge, E. Shi, and W. Dorland. Discontinuous Galerkin algorithms for fully kinetic plasmas. *Journal of Computational Physics*, 353:110–147, 2018.

[16] James Juno, Ammar Hakim, and Jason M. TenBarge. A parallel-kinetic-perpendicular-moment model for magnetised plasmas. *Journal of Plasma Physics*, 91(5):E129, 2025.

[17] H. J. Kull. Theory of the Rayleigh–Taylor instability. *Physics Reports*, 206(5):197–325, 1991.

[18] A. Lenard and Ira B. Bernstein. Plasma oscillations with diffusion in velocity space. *Physical Review*, 112(5):1456–1459, 1958.

[19] John Rodman, Petr Cagas, Ammar Hakim, and Bhuvana Srinivasan. Kinetic interpretation of the classical Rayleigh–Taylor instability. *Physical Review E*, 105(6):065209, 2022.

[20] John Rodman, James Juno, and Bhuvana Srinivasan. The effect of spatially varying collision frequency on the development of the rayleigh–taylor instability. *The Astrophysical Journal*, 965:173, 2024.

[21] Yang Song and Bhuvana Srinivasan. A survey of the effects of magnetic fields, resistivity, viscosity and thermal conduction on the Rayleigh–Taylor instability. *Radiation Effects and Defects in Solids*, 175(11–12):1009–1014, 2020.

[22] Bhuvana Srinivasan, Guy Dimonte, and Xian-Zhu Tang. Magnetic field generation in rayleigh–taylor unstable inertial confinement fusion plasmas. *Physical Review Letters*, 108(16):165002, 2012.

[23] Bhuvana Srinivasan and Xian-Zhu Tang. Mechanism for magnetic field generation and growth in rayleigh–taylor unstable inertial confinement fusion plasmas. *Physics of Plasmas*, 19(8):082703, 2012.

[24] Bhuvana Srinivasan and Xian-Zhu Tang. The mitigating effect of magnetic fields on rayleigh–taylor unstable inertial confinement fusion plasmas. *Physics of Plasmas*, 20(5):052109, 2013.

[25] Bhuvana Srinivasan and Xian-Zhu Tang. Mitigating hydrodynamic mix at the gas–ice interface with a combination of magnetic, ablative, and viscous stabilization. *Europhysics Letters*, 107(6):65001, 2014.

[26] Takayuki Umeda and Yasutaka Wada. Secondary instabilities in the collisionless Rayleigh–Taylor instability: Full kinetic simulation. *Physics of Plasmas*, 23(11):112117, 2016.

[27] Takayuki Umeda and Yasutaka Wada. Non-MHD effects in the nonlinear development of the MHD-scale Rayleigh–Taylor instability. *Physics of Plasmas*, 24(7):072307, 2017.

[28] Liang Wang, K. Germaschewski, Ammar H. Hakim, and A. Bhattacharjee. Comparison of multi-fluid moment models with particle-in-cell simulations of collisionless magnetic reconnection. *Physics of Plasmas*, 22(1):012108, 2015.

Appendix A

INPUT FILES AND CODE AVAILABILITY

The simulation input files, post-processing scripts used in this thesis are available in the following GitHub repository:

`https://github.com/tej3000/MS-Thesis/tree/main`

The repository contains the numerical input files for the kinetic and PKPM simulations, the scripts used to extract interface amplitudes and growth rates, and the code used to generate the figures presented in this thesis.