FLOW PHYSICS OF RADIATIVELY HEATED PARTICLE-LADEN CHANNEL FLOW AND SIMULATION METHODS FOR SHOCK -DRIVEN PROBLEMS IN MATERIALS WITH STRENGTH

By
Jacob Roy West and Sanjiva Lele

Report No. TF-192

Flow Physics and Computational Engineering Group
Department of Mechanical Engineering
Stanford University
Stanford, CA 94305

March 2023

Flow Physics of Radiatively Heated Particle-Laden Channel Flow and Simulation Methods for Shock-Driven Problems in Materials with Strength

By

Jacob Roy West and Sanjiva Lele

Report No. TF-192

Flow Physics and Computational Engineering Group
Department of Mechanical Engineering
Stanford University
Stanford, CA 94305

March 2023

Abstract

Simulations of multiphase flows are used to understand and predict a wide variety of natural and industrial phenomena. This thesis concerns simulations of two different multiphase flows.

The first part of this thesis uses high fidelity simulations to study particle-laden turbulent channel flow, subjected to strong radiation heating. Chapter 2 describes the physics of isothermal particle-laden channel flow, characterizing the effects of particles on turbulence across wide parametric variation in Stokes number and mass loading. Ideas from other branches of turbulence study, in particular compressible turbulence and rough-wall turbulence, are applied to make progress toward a model for two-way coupled particle-laden channel flow. Chapter 3 considers the dynamics of particle clouds in isothermal particle-laden channel flow using new tessellation based approaches to characterize particle clustering, rotation, and swirling motions. The dependence of the particle dynamics is characterized as function of Stokes number and mass loading, but also flow region, and comparisons with previous results from homogeneous isotropic turbulence are made in the logarithmic layer. Finally, the effect of strong radiation heating is shown on particle-laden channel flow in chapter 4. The strong heat transfer drives expansion and acceleration, and the flow tends towards re-laminarization from acceleration and viscosity increase.

The second part of this thesis describes the development of an Eulerian finite difference method for simulating multi-phase mixtures of elastic-plastic materials subjected to shock waves and undergoing strain hardening. The new method uses localized artificial diffusivity (LAD) to regularize discontinuities, including discontinuities in strain, which are important in shocks in solids. New LAD terms are added to stabilize the kinematic evolution equations in strong shear deformation. The new method is tested on a variety of 1D and 2D problems to demonstrate its convergence properties, ability to simulate impacts, and shock capturing ability. The 2D problems considered are a Taylor impact with realistic equation of state and plasticity models, and a solid-solid Richtmyer-Meshkov instability, which would be challenging for more traditional Lagrangian methods. Lastly, three interface sharpening techniques for multiphase flows are extended to elastic-plastic materials, and their performance is evaluated on a variety of test problems. The methods are comparable when simulations are well-resolved, but show different deficiencies when flow topologies become under-resolved.

Acknowledgments

THIS PAGE DELIBERATELY LEFT BLANK

Contents

List of Tables

List of Figures

xiii

Chapter 1

Introduction: Flow Physics of Radiatively Heated Particle-Laden Channel Flow

Multiphase flow is a branch of fluid dynamics concerned with the motion of a mixture of two or more materials in which one or more of the constituents moves in a fluid-like manner. Some examples include dust storms in the atmospheric boundary layer (solid-gas mixture), spray atomization of liquid fuel in an air environment (liquid-gas), and emulsions of oil and water (liquid-liquid). The study of multiphase flows is commonly divided into dispersed and separated flows [32]. In separated multiphase flows, there are distinct regions of each continuous fluid or solid phase, and the interfaces among them may change in shape over time. Examples include bubbly flows generated by breaking ocean waves and cavitation on boat propellers. Dispersed multiphase flows comprise a carrier fluid phase containing discrete particles, droplets, or bubbles, which are transported by the carrier phase, but may also exert forces upon it. Examples of dispersed particle-laden flows include soot production and transport in diesel engines, and sediment transport in rivers, as shown in figure 1.1.

Accurate simulations of multiphase flows are important because they enable engineering predictions of many of the previously mentioned real-world applications, but there are significant difficulties in simulating these flows. In separated flows, the difficulty arises from the different equations describing the carrier fluid (Navier-Stokes equations) and the suspended phase (ordinary differential equations describing the trajectories of particles). In separated flows, the Navier-Stokes equations are typically solved throughout both media, but if the media are very dissimilar, then there are sharp gradients in properties such as density and viscosity, as well as interfacial physics in thin regions, such as surface tension, which are difficult to resolve. This dissertation concerns the simulation of two very different multiphase flows. Chapters 2-4 describe how simulations of particle-laden channel flow

FIGURE 1.1: Two examples of dispersed particle-laden flows: soot plumes from a diesel engine (left, EPA, Public domain, via Wikimedia Commons) and sediment transport at the outlet of the Tiber river (right, Copernicus Sentinel data (2019), processed by ESA,CC BY-SA 3.0 IGO).

with intense radiation heat transfer are used to better understand the physics of this system. The second part, chapters 5-7, describes the development of numerical methods for simulating mixtures of elastic-plastic materials subjected to shocks.

When simulating particle-laden turbulent flows, there are two primarily used methods: particle-resolved, in which every flow length scale is resolved, including the wakes and boundary layers around individual particles; and point-particle methods, in which individual particles are tracked, but only interact with the fluid through forces they exert on it [31]. The particle-resolved method can be more accurate, by nature of its fewer approximations, and increasing computational resources are making it possible to use in more and more scenarios [44, 45]. In this work, because of the number of particles used and their small size, the point-particle method is used. We leverage recent advances [87, 53] in point-particle methods to ensure the accuracy of these simulations while keeping the computational cost affordable.

The simulations of particle-laden flow in this dissertation are inspired by a particular application, known as a particle solar receiver, which was studied as the subject of the PSAAP-II Center at Stanford. In this device, originally proposed by Hunt [90] and Abdelrahman et al. [1], concentrated solar radiation is absorbed by solid particles, falling through or suspended in a transparent gas, to heat a power generation cycle in a power plant. The benefit of such a design would be increased surface area of the solar receiver and ability to achieve higher temperatures, as compared to traditional
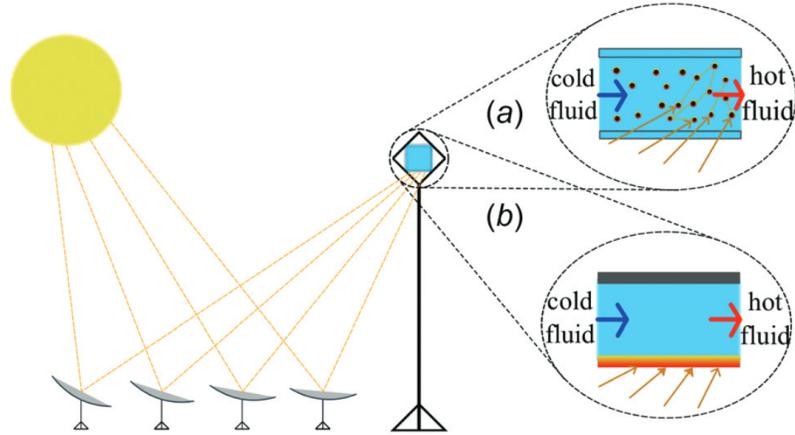
FIGURE 1.2: Schematic of concentrating solar receiver, taken from Pouransari and Mani [164]. (a) Shows the particle-laden flow receiver concept, and (b) shows the more traditional receiver, with an opaque working fluid.

working fluids in concentrating solar, such as molten salts [86]. A schematic of this device is shown in figure 1.2.

This PSAAP-II center conducted a variety of experimental and computational studies to understand the physics of this device [16, 55, 164, 165, 64]. However, because of constraints on the experimental apparatus, the majority of studies in the center utilized low mass loading and radiation intensity, with little variation of particle properties. The overall goal of the study presented in chapters 2-4 is to better understand the flow physics of particle-laden wall-bounded flows in a mass loading and radiation intensity regime closer to the proposed device than could be achieved in the laboratory setting. This is accomplished by first studying in detail particle-laden channel flow without heating. In chapter 2, we thoroughly characterize the modulation of turbulence by particles in channel flow, and assess how the transfer of momentum and turbulent kinetic energy are affected by particles, considering both the mass loading and particle inertia as parametric variation. In doing so, we leverage ideas from the study of compressible wall turbulence to relate the modifications to the velocity profile back to the single phrase, incompressible channel flow, about which more is known [201, 198, 158].

In chapter 3, we pivot to the particle phase in the unheated flow, using new techniques based on tessellation introduced by Oujia et al. [151, 153] to study the clustering, rotation, and swirling motion of particle clouds. Because of the two-way coupling which occurs at large mass loading, these motions affect the fluid-phase dynamics, especially at large mass loading, and so this study shines new light onto the turbulence modulation observed in the fluid phase. Finally, in chapter 4 we conclude this part of the thesis by presenting simulation results of particle-laden channel flow at large mass loading with radiation intensity closer to the intended applications, with roughly 1500

times the intensity of solar radiation. We observe a sharp adjustment after the radiation is applied, after which the flow reaches a new equilibrium between the acceleration, and the increase of the wall shear stress. The baseline flow established in the unheated simulations remains very similar in terms of its turbulent structure, but that a strong degree of heating, expansion, and acceleration is superimposed on it. This results in attenuation of the turbulence through the acceleration and also through viscosity increase with temperature. The modifications to the flow, in terms of relevant engineering quantities such as temperature rise, can be predicted reasonably well with a simplified one-dimensional model, even at these more realistic conditions.

Chapter 2

Unheated particle-laden channel flow

2.1 Introduction

Particle-laden turbulent flows over walls occur in many natural and engineered systems. Examples include sediment transport in rivers, dust storms in the atmospheric boundary layer, and chemical processes in industrial riser reactors. A dispersed particle phase suspended in a gas may also be used as a novel working fluid in high-temperature solar energy applications, in which solid particles absorb concentrated solar radiation [134, 183, 90, 1]. In order to have a respectable thermal efficiency, such a device would need a sufficiently high optical depth or volume fraction to absorb most of the incident radiation in the particle phase. With solid particles, this naturally leads to large mass loadings. Therefore, understanding the flow physics of particle-laden wall-bounded turbulence at high mass loading is important to designing and predicting performance of these devices. In particular, we focus on heavy particles, i.e. with a large density relative to the carrier fluid, which would be typical for solid particles in a gas. The key motivating questions in this study are (1) How are the transport of mass, momentum, and energy in wall-bounded particle-laden flows altered as a function of mass loading and Stokes number, and (2) Can scaling and transform results from the compressible flow literature be leveraged to relate the behavior of the multiphase mixture to incompressible wall-bounded flows?

While single-phase incompressible channel flow depends only on the Reynolds number, particle-laden flows exhibit many phenomena depending on the Stokes number, mass loading ratio, volume fraction, and density ratio of particle to fluid, to name a few important parameters [13, 31]. As a result, there have been many experimental and computational studies exploring the effects of these parameters in pipe and channel flows. We do not review these studies exhaustively here, but focus on

flows with mass loading ratio ϕ, defined as the ratio of the mass of the dispersed and carrier phases, is $\mathcal{O}(1)$ while the volume fraction is still dilute (less than 0.001). In this regime, both the mean and fluctuating flow quantities are significantly affected by inter-phase coupling. In the following literature review, we differentiate studies based on the average mass loading ϕ_0, the friction Stokes number $St^+ = \tau_p/\tau_{visc}$, and the friction Reynolds number $Re_\tau = u_\tau \delta/\nu$, where τ_p is the particle relaxation time, τ_{visc} is the viscous time scale associated with the wall shear stress, u_τ is the friction velocity, δ is the channel half height, and ν is the kinematic viscosity.

In an experimental study by Kulick et al. [113], there was no significant modification to mean fluid velocity profiles observed for mass loadings from 0.02 to 0.4, but they note the difficulty in obtaining accurate measurements for flows with higher mass loadings. They did observe reductions in fluid velocity fluctuation intensities as mass loading increases across this range, which has since been documented in numerous computational studies.

In an early computational study of particle-laden channel flow at low Reynolds number ($Re_\tau = 125$), Li et al. [117]. evaluated the effects inter-particle collisions, density ratio, and Stokes number at a range of mass loadings from 0-2. They found that collisions strongly affected the rate of wall accumulation, and that particles tended to increase the anisotropy of fluid phase Reynolds stresses. The degree of turbulence modulation was strongly affected by the mass loading ratio, and more subtly affected by Stokes number and density ratio. Lee and Lee [114] thoroughly explored the effect of Stokes number at mass loading of 0.3, finding that particles with very low inertia could enhance the fluid turbulence, while larger particles attenuate it, and that maximum preferential concentration of particles occurred around $St^+ \approx 5$. In a series of works, Capecelatro and co-authors examined the effect of mass loading and Stokes number in a vertical channel, considering flow in the same direction as, and opposite to gravity [33, 34, 35]. They found that for mass loadings much greater than 1, the dominant turbulence production method in riser flows shifted from shear production to "drag production" driven by the force of gravity on particle clusters. In the transition between these two modes, the flow becomes nearly laminar as shear production is suppressed.

While the work of Capecelatro et al. [35] considered a wider range of mass loadings than presented here, we provide additional intermediate cases where mass loading is $\mathcal{O}(1)$ to resolve the changes in turbulence in this range. Further, since we neglect gravity, the present study can show the degree to which the particle inertia also plays a role in re-laminarizing the flow as mass loading is increased, independent of body forces. Additionally, the present study considers a wider range of Stokes numbers, including much larger Stokes numbers.

Vreman [203] studied particle-laden vertical pipe flow over a mass loading range of 0.11 to 30, in which the flow was aligned with gravity. This study was conducted for $Re_\tau = 140$ and large Stokes number, and found that turbulence is significantly suppressed with larger mass loading. Using integrated Reynolds stress budgets, Vreman argues that reductions in turbulence intensity in particle-laden flow can be understood as occurring by similar mechanisms as in compressible

flows, although the physics driving the variable inertia is different [203]. In this paper and later work on channel flow [204], the importance of collisions for accurate prediction of wall-bounded flows is demonstrated. Even when the volume fraction is relatively dilute, collisions are important in establishing the equilibrium concentration profile in a particle-laden flow. The importance of collisions was also noted by Yamamoto et al. [212], who used large-eddy simulation to study particle-laden channel flows with mass loading between 0.01 and 1.0 at $Re_\tau = 644$. They found that particle RMS velocities and particle number density distributions are are strongly influenced by collisions even when the volume fraction is very dilute, i.e. $\mathcal{O}(10^{-4})$. Accounting for collisions improved agreement with experiment for particle velocity fluctuations and the particle distribution. At mass loading of 1, they also observe a substantial increase in the mean fluid velocity, although this is driven by gravity.

Recently, Costa et al. [45] studied near-wall turbulence modulation with small particles up to mass loadings of 0.5, with the goal of reconciling conflicting reports about two-way coupled channel flows, such as observations of drag increase and decrease. They used particle-resolved simulation to provide high fidelity validation that is a useful reference for the study here, which uses the point particle technique to examine even higher mass loadings. Yu et al. [214] also performed particle-resolved simulations, with larger particles at mass loadings up to 2.36 in an upward channel flow. They used their study to more accurately characterize the criteria for whether particles attenuate or enhance fluid turbulence, finding that this depends on particle Reynolds number, density ratio, particle size, and volume fraction. Xia et al. [211] studied the downward flow configuration with particle-resolved simulations, characterizing drag increases due to particle loading, as well as criteria for turbulence attenuation/enhancement.

Zhao et al. [218] studied energy transfer between fluid and particle phases in two-way coupled channel flow without collisions. They performed simulations up to mass loading of 2.0, at three different Stokes numbers, and observed drag reduction as mass loading was increased. They also described how turbulence kinetic energy is augmented or attenuated by particles, depending on the flow parameters. More recently, Pan et al. [155] also studied turbulent and kinetic energy budgets using two-way coupled simulations at mass loading of 1.0. Both of these works describe how particles transfer energy to the fluid near the wall and receive energy from the flow in the channel center. Zhou et al. [219] presented detailed TKE budgets of at several mass loadings up to 1.0, also without inter-particle collisions. They showed the frequently reported attenuation of wall-normal velocity fluctuations and enhancement of streamwise velocity fluctuations, but noted non-monotonic behavior with mass loading in the strength of streamwise fluctuations.

In all of the above works, the Reynolds number is fairly low, with high values in the literature being the direct numerical simulations of Capecelatro et al. [35] at $Re_\tau = 300$, although higher values have been done in the large-eddy simulation of Yamamoto et al. [212] at $Re_\tau = 644$ and the experimental comparison with duct flow at $Re_\tau = 570$ in Esmaily et al. [55], although in that study

the highest mass loading was 0.12. The lack of simulations at high mass loading and high Reynolds number is understandable given the expense involved in computing inter-particle collisions for large numbers of particles while time-stepping is limited by the fine scales of fluid turbulence. Further, the point particle approximation becomes worse as the Kolmogorov scale becomes small relative to particle diameter. Large eddy simulation could reduce these costs, but the vast majority of research on sub-grid scale modeling for particle-laden flows has focused on the one-way coupled regime [124], with more development work needed for strongly coupled flows.

In the literature on particle-laden flows in channels and pipes reviewed here, the effect of the presence of the particle phase on mean velocities and turbulent statistics has been well-documented. Relative to single-phase incompressible channel flow, the changes to the mean streamwise velocity profile include changes in near-wall slope, log layer slope, and log layer offset. These types of changes are reminiscent of the types of changes observed in compressible or variable-property wall-bounded flows. Efforts to relate these modifications via scaling or transformations to incompressible flow go back to Van Driest [201]. Zhang et al. [216] improved upon the Van Driest transform by invoking the balance between shear production and viscous dissipation, and by using mixing length arguments. This improved over Van Driest for adiabatic boundary layers, but not other settings. A different approach, which better scales viscous stresses and properly accounts for near-wall gradients in fluid properties, was derived independently by Trettel and Larsson [198] and by Patel et al. [158]. Subsequent papers by Patel, Pecnik, and co-workers demonstrate how the improved transformation opens up possibilities for better RANS models in compressible and variable property flows [159, 150]. A further refined transform, which synthesizes the aforementioned approached, and has wider applicability in terms of Mach number and wall boundary condition, was recently proposed by Griffin et al. [78]. Since it is known that the presence of particles can alter both the inertia and viscosity of a multiphase mixture [111], this suggests that a transform approach might be applicable to modulation of particle-laden flows. Such an approach has in fact been successfully used for particle modified viscosity by Costa et al. [43], who found that it was also necessary to account for the effects of particle slip near the wall, in a manner reminiscent of an offset in a rough-wall flow. However, the dense multiphase mixtures they examine do not feature strong gradients in particle-induced viscosity, and feature only neutrally buoyant particles, so that variable density effects are not present.

The emphasis of the present work is to study a wider range of particle mass loadings and Stokes numbers using recent developments in point-particle approaches [87, 53] to correct for the undisturbed velocity, and characterize the changes brought on by particle coupling in terms of scaling of mean velocities, and to connect this scaling to compressible flow transformations and rough wall flows. The remainder of this paper is organized as follows: Section 2.2 describes the computational method and suite of simulations used. Section 2.3 describes the effects of mass loading and Stokes number on particle-laden channel flow. In section 2.4, the prospect of a transform for relating this

data to incompressible flow is investigated. Section 2.5 presents a summary and conclusions.

2.2 Methods

The four-way coupled particle-laden channel flow simulations were performed using the Soleil-MPI code, which is briefly described here; more details on this code can be found in [55]. The Navier-Stokes equations for incompressible flow with particle two-way coupling terms are solved as

$$\nabla \cdot \boldsymbol{u} = 0 \tag{2.1}$$

$$\frac{\partial (\rho \boldsymbol{u})}{\partial t} + \nabla \cdot (\rho \boldsymbol{u} \otimes \boldsymbol{u}) = -\nabla p + \mu \nabla^2 \boldsymbol{u} + \sum_{m=1}^{N_p} \boldsymbol{f}^{(m)} \delta_{\text{Dirac}}(\boldsymbol{x} - \boldsymbol{x}_p{}^{(m)}), \tag{2.2}$$

where $\boldsymbol{u}$ is the fluid velocity, ρ is the (constant) fluid density, p is the hydrodynamic pressure and μ is the dynamic viscosity. The last term is the summation of the two-way coupling force contribution from each particle $\boldsymbol{f}^{(m)}$, over all particles N_p, based on the Dirac-delta function $\delta_{\text{Dirac}}(\cdot)$, which has units of inverse volume, and whose argument is the distance between a position in the fluid $\boldsymbol{x}$ and a given particle $\boldsymbol{x}_p{}^{(m)}$. Because the focus in this work is on flows with a small particle volume fraction (less than 1.5×10^{-4}), the volume displacement effect, which should be negligible in this regime, has been ignored [154]. The particle drag force is the drag correlation for finite Reynolds numbers from [177], which is given by

$$\boldsymbol{f}^{(m)} = 3\pi \mu D_p \left(1 + 0.15 Re_p^{0.687}\right) \left(\boldsymbol{v}_p^{(m)} - \widetilde{\boldsymbol{u}}^{(m)}\right), \tag{2.3}$$

where D_p is the particle diameter, $\boldsymbol{v}_p^{(m)}$ is the velocity of particle m, $\widetilde{\boldsymbol{u}}^{(m)}$ is the undisturbed fluid velocity at the particle location, and Re_p is the particle Reynolds number. The Basset history force term in the Maxey-Riley equation [130] is neglected on account of the larger density ratio (~ 7400) used in this work.

Soleil-MPI solves the above equations using a finite-volume method with second-order explicit spatial derivatives. Variables are stored at staggered locations: velocities at cell faces and pressure and cell centers. At each timestep, a pressure Poisson equation is solved to achieve a divergence-free velocity field, using a geometric multigrid method designed for non-uniform grids and variable density flows [54]. The two-way-coupling terms in the momentum equations are obtained by distributing the force associated with each particle to the surrounding grid points using tri-linear interpolation. It is well-known that the undisturbed fluid velocity in the Stokes drag equation is not known in a two-way coupled simulation, which can lead to underestimates of the drag force [87]. Here we use the correction scheme for anisotropic grids proposed by [53] to account for particle disturbances on the fluid velocity.

Each particle trajectory is computed according to the equations of motion for a single particle

of mass m_p, neglecting gravity, as

$$\dot{\boldsymbol{x}}_p^{(m)} = \boldsymbol{v}_p^{(m)}, \tag{2.4}$$

$$m_p \dot{\boldsymbol{v}}_p^{(m)} = -\boldsymbol{f}^{(m)}. \tag{2.5}$$

Recent works have indicated that gravity can have a significant effect on particle distribution and dynamics, depending on the Stokes number and other flow parameters [115, 30], but we neglect this effect here for simplicity. Time advancement for both the Eulerian fluid variables and the Lagrangian point particles is done with a second-order Runge-Kutta method. Collisions with walls and other particles are accounted for assuming perfectly elastic collisions according to the hard sphere model [49], using the algorithm described in Rydquist and Esmaily [171] to detect collisions. The cost of this method scales linearly with the number of particles, which prevents the collisions from overwhelming the computational cost.

A suite of particle-laden channel flow cases, described in table 2.1, are used in this work to assess the role of mass loading and Stokes number. Mass loading is defined as the ratio of particle mass to fluid mass: $\phi = nm_p/\rho$, where n is the particle number density. Each case is described by the mass loading averaged over the domain, ϕ_0, but ϕ varies in time and space for each flow due to preferential concentration and turbophoresis. Due to turbophoresis, the mass loading decreases with wall distance, as indicated by the average mass loadings for each flow region in table 2.1. The Stokes number primarily used in this work is the friction Stokes number, $St^+ = \tau_p/\tau_{visc}$, which is the ratio of the particle relaxation time scale $\tau_p = \rho_p D_p^2/(18\mu)$ to the viscous time scale associated with the wall shear stress $\tau_{visc} = \delta/\sqrt{\tau_w/\rho}$. The channel half-height is δ, and τ_w is the wall shear stress. Table 2.1 also includes the Kolmogorov Stokes number $St_\eta^{\log}$, which is defined based on the dissipation rate ϵ of turbulent kinetic energy averaged over the range $50 < y^+ \leq Re_\tau$, i.e., $St_\eta^{\log} = \tau_p\sqrt{\rho\epsilon/\mu}$. This represents an average Stokes number in the logarithmic region, to facilitate comparison with results from homogeneous isotropic turbulence. Here, y^+ is the wall-normal coordinate normalized by the viscous length scale $l_{visc} = \mu/\sqrt{\rho\tau_w}$. The cases in which St^+ is varied all have $\phi_0 = 1$, at which inter-phase coupling is strong, and the cases in which ϕ_0 is varied all have $St^+ \approx 7$, at which a strong degree of preferential concentration is expected based on the simulations of Lee and Lee [114].

The fluid and particle properties are based on the benchmark experiments and computations described in [55], with a channel geometry instead of a duct. These experiments used small nickel particles in a turbulent air flow. Accordingly, in our simulations, the carrier fluid is air, with $\rho = 1.2$ kg/m^3 and $\mu = 1.87 \times 10^{-5}$ Pa-s, and the particles are monodisperse spheres with diameter $D_p = 11.5$ μm. The average mass loading ϕ_0 is varied by changing the number of particles, using a constant particle density of $\rho_p = 8900$ kg/m^3 (the density of nickel). The Stokes number St^+ is varied by changing the particle density and number of particles at a constant ϕ_0. Across all simulation cases, the fluid mass flux is held constant. As N_p and ρ_p are varied, Re_τ varies slightly

	Re_τ	N_p	ρ_p/ρ	ϕ_0	ϕ_{visc}	ϕ_{buf}	ϕ_{log}	St^+	$St^{\log}_\eta$
Flow 0	229	0	N/A	0	0	0	0	N/A	N/A
Flow 1	225	1.43×10^7	7.42×10^3	0.1	1.37	0.13	0.06	6.90	0.69
Flow 2	227	5.70×10^7	7.42×10^3	0.4	2.36	0.59	0.32	7.01	0.59
Flow 3	234	1.43×10^8	7.42×10^3	1.0	3.54	1.41	0.88	7.45	0.54
Flow 4	258	2.85×10^8	7.42×10^3	2.0	5.16	2.90	1.82	9.04	0.53
Flow 5	273	4.28×10^8	7.42×10^3	3.0	5.66	4.23	2.79	10.1	0.53
Flow 6	263	1.43×10^9	7.42×10^2	1.0	1.52	1.12	0.98	1.02	0.075
Flow 7	273	4.51×10^8	2.35×10^3	1.0	2.63	1.31	0.93	2.98	0.21
Flow 8	213	4.51×10^7	2.35×10^4	1.0	2.79	1.58	0.85	19.6	1.42
Flow 9	210	1.43×10^7	7.42×10^4	1.0	2.10	1.47	0.89	60.0	3.66

TABLE 2.1: Parameters of the channel flow data: the friction Reynolds number Re_τ, the number of particles N_p, the particle to fluid density ratio ρ_p/ρ, the overall mass loading ϕ_0, the average mass loading in the viscous sublayer ϕ_{visc}, the average mass loading in the buffer layer ϕ_{buf}, the average mass loading in the logarithmic layer ϕ_{log}, the friction Stokes number St^+, and the Kolmogorov Stokes number based on the dissipation rate averaged over the logarithmic layer $St^{\log}_\eta$.

from the no-particle value. This indicates either drag reduction (Re_τ increase) or drag enhancement (Re_τ decrease). Small amounts of drag reduction due to two-way coupling have been observed in some other studies, as described by Zhao et al. [218]. The small changes in Re_τ contribute to small changes in St^+, as indicated in table 2.1.

All simulation cases use the same flow domain and boundary conditions. The channel dimensions are $4\pi\delta \times 2\delta \times (4/3)\pi\delta$ (where $\delta = 2$cm), with periodic boundary conditions in x and z, and smooth, no-slip walls in y. These domain dimensions are sufficient to allow de-correlation of velocity fluctuations in the streamwise and spanwise directions for the flow without particles [110, 139]. It has been shown that even in one-way coupled simulations, the particle concentration is affected by domain size, and that larger domains are required for de-correlation as Stokes number is increased [173]. The results presented in chapter 3 suggest that the domain used is sufficient for low Stokes number, but the domain may be too small for large Stokes numbers (see figure 3.23, for example). A grid resolution of $[N_x, N_y, N_z] = [280, 140, 140]$ was used to resolve the fluid turbulence at reasonable computational cost. The computational grid is stretched in the wall-normal direction using the hyperbolic tangent stretching function from Bose [28], Eqn. 2.11. The grid is uniform in homogeneous directions. For the single-phase baseline case, there are 23 grid points below $y^+ = 10$. The centerline grid resolution is $[\Delta x^+, \Delta y^+, \Delta z^+] = [10.5, 9.1, 7.0]$. In appendix 2.6, the effect of one level of grid refinement to $[N_x, N_y, N_z] = [560, 280, 280]$ on various statistics for Flows 0 and 3 is reported.

A constant time step was used for each simulation case, based on the viscous stability condition, such that for all cases $\Delta t^+ < 9 \times 10^{-3}$. For the statistics reported in the following sections, all simulation cases were run until a steady state was reached. Then, simulations were continued for a minimum of 5 flow-through times. Eulerian statistics of time-averaged quantities were obtained by
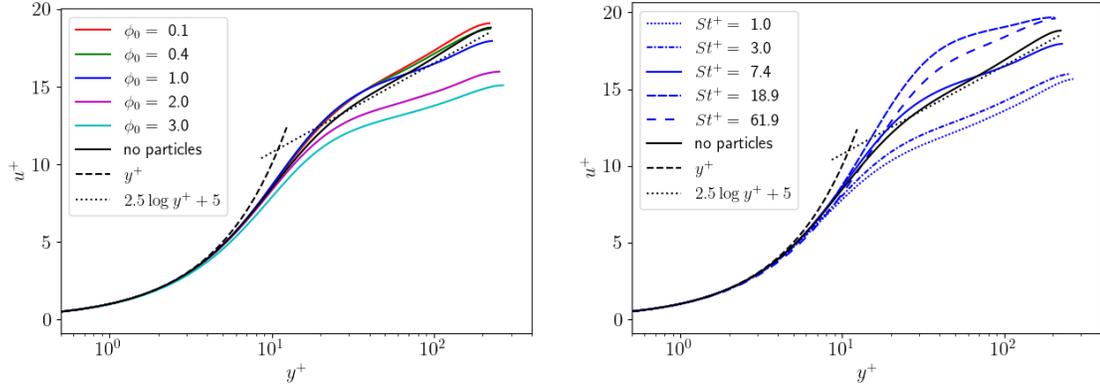
FIGURE 2.1: Mean fluid velocities for simulation cases spanning variation in ϕ_0 (left) and St^+ (right). The near-wall viscous scaling and incompressible logarithmic profile are plotted for reference.

averaging continuously throughout this period, as well as over homogeneous directions.

2.3   Turbulence modulation by inertial particles

In this section, we present mean and turbulent statistics from the suite of channel flow cases, exploring the influence of mass loading and Stokes number. In all cases, it is important to remember that the fluid mass flux is held constant, which means that both the total mass flux and wall shear stress vary among the cases. In this section, viscous scales will be used as the normalization, and non-dimensional quantities are indicated with a superscript "$+$". A different approach to scaling and normalization will be explored in section 2.4.

Figure 2.1 shows the mean streamwise fluid velocities across the variation of mass loading and Stokes number. Depending on the flow parameters, the centerline velocity may be less or greater than the value for the incompressible equivalent, which is an indicator of drag increase or decrease, respectively. There is also a clear dependence of the log layer slope on the overall mass loading ϕ_0, with larger mass loadings showing a shallower slope. The variation in St^+ does not cause a similar degree of change in the log layer slope, but instead changes the offset of the log layer. This is reminiscent of a roughness effect, which will be explored further in section 2.4. Across all the data, streamwise velocity profiles remain reasonably collapsed up to $y^+ \approx 10$. This implies that the traditional viscous scales are appropriate through the viscous sublayer and into the buffer layer, but that the buffer and log layer dynamics are altered by the presence of particles.

Next, the local particle mass loading is shown in figure 2.2. Across most cases, we observe a thin layer at the wall (below $y^+ \approx 2$ where the local mass loading is increasing with wall-normal distance. Above this, the mass loading decreases toward the centerline value, which is slightly less than the nominal value. This decreasing behavior due to turbophoresis has been widely observed in
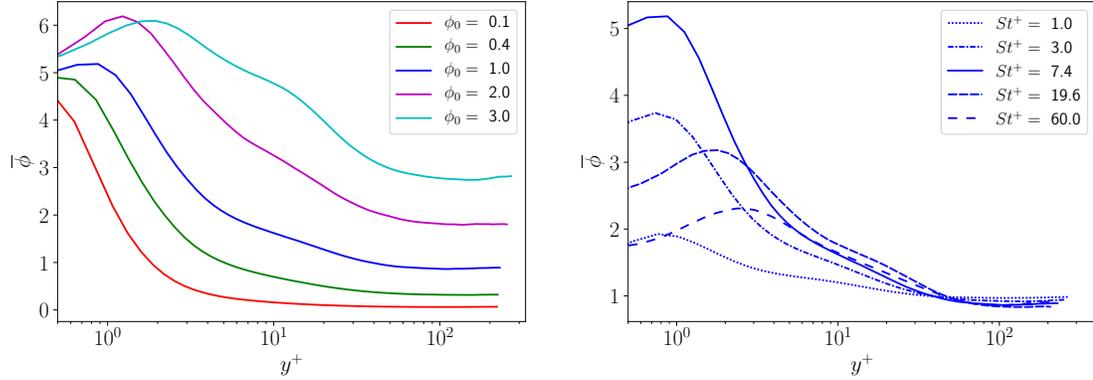
FIGURE 2.2: Local average mass loading for simulation cases spanning variation in ϕ_0 (left) and St^+ (right).

turbulent flows, and is expected on a theoretical basis as well [103]. The near-wall increase suggests that collisions play an increased role in that region. In the ϕ_0 variation, the particle attributes are kept constant, so it is not surprising that a similar equilibrium value of local mass loading is achieved, despite the differences in centerline particle concentration. As mass loading is increased, the peak in mass loading reaches up higher into the inner layer, and the particles are more evenly distributed across the channel. When St^+ is varied at constant ϕ_0, a similar centerline value of ϕ is reached, but the near-wall values vary non-monotonically, with an intermediate value ($St^+ = 7.4$) showing the highest tendency toward wall accumulation. The peak location is generally observed to increase with St^+.

To characterize the effect of particles on fluid turbulence, the fluid velocity fluctuations are shown in figure 2.3. In agreement with previous studies, these show that wall-normal and spanwise velocity fluctuations are attenuated as both ϕ_0 and St^+ are increased. Streamwise fluctuations, however are less severely affected, with some cases even showing enhancement with respect to the unladen flow. Together, these observations lead to the conclusion that when particles are added to a channel flow, the turbulence becomes less isotropic, and more dominated by streamwise fluctuations. This is due to an attenuation of the pressure-strain term in the TKE budget, which prevents transfer of energy from the streamwise to other directions [117].

The effect of particles on anisotropy is shown more directly in figure 2.4, which shows the barycentric map of the invariants of the Reynolds stress anisotropy tensor [15]. This shows that as mass loading or Stokes number are increased, the Reynolds stress moves closer to the one-component limit, reflecting the dominance of streamwise velocity fluctuations over spanwise and wall-normal fluctuations.
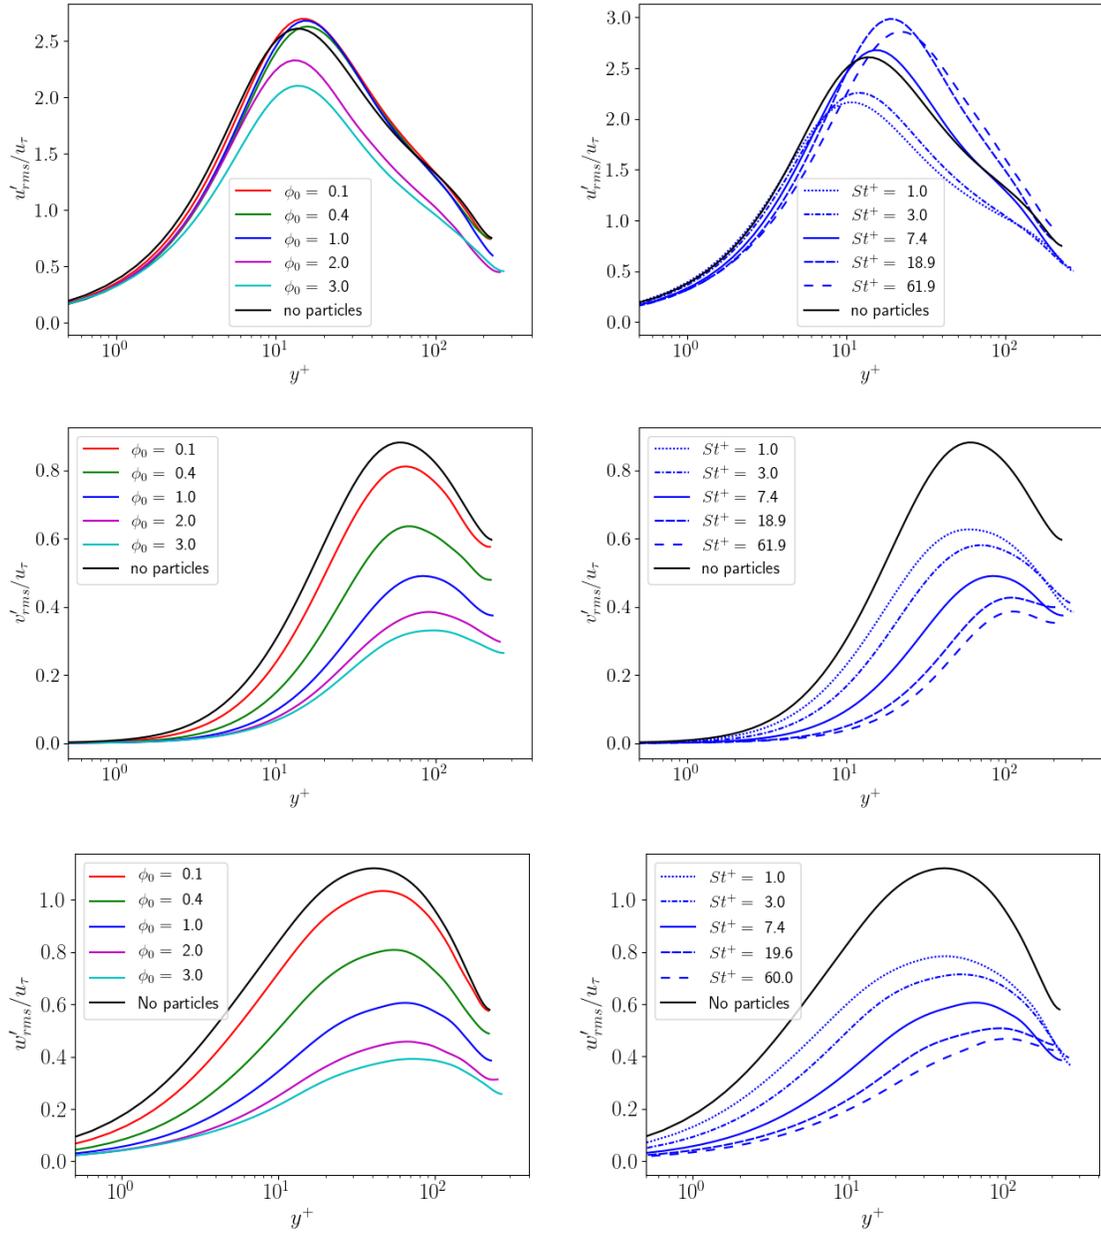
FIGURE 2.3: Comparison of RMS fluid velocity fluctuations in the streamwise (top row), wall-normal (middle row), and spanwise directions (bottom row), spanning variation in ϕ_0 (left) and St^+ (right).
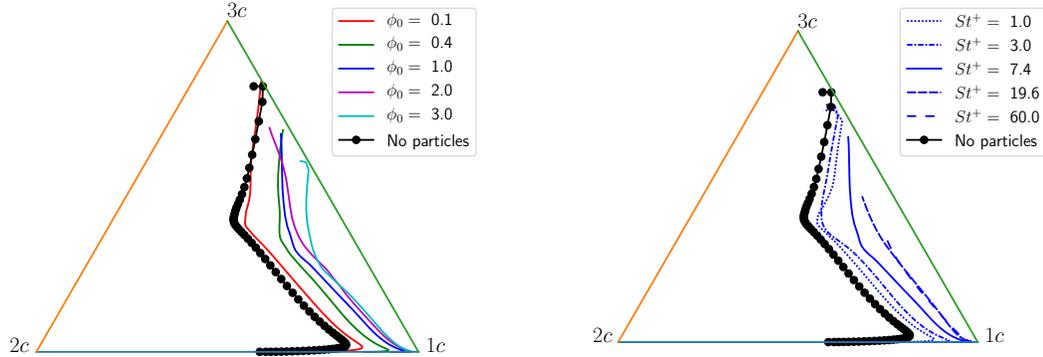
FIGURE 2.4: Barycentric map of the invariants of the Reynolds stress anisotropy tensor for cases spanning variation in ϕ_0 (left) and St^+ (right). "1c", "2c", and "3c" denote the one-, two-, and three- component limiting states.

2.3.1   Momentum budgets

The attenuation of wall-normal velocity fluctuations as a function of St^+ and ϕ_0 shown in figure 2.3 suggest significant changes to the Reynolds shear stress, and thus to the stress balance in the channel. The stress balance for a fully developed particle-laden channel flow is given by

$$(1 - \frac{y}{\delta}) = \frac{1}{\tau_w} \left(\mu \frac{d\overline{u}}{dy} - \rho \overline{u'v'} + \overline{F}_x \right), \tag{2.6}$$

where F_x is the streamwise component of the sum of the particle feedback forces, i.e. $F_x = \sum_{m=1}^{N_p} f_x^{(m)} \delta_{\text{Dirac}}(\boldsymbol{x} - \boldsymbol{x}_p^{(m)})$. The Reynolds decomposition $f = \overline{f} + f'$ has been used to decompose quantities into their mean $\overline{}$ and fluctuating $'$ components, respectively. This term contributes to the stress balance between the viscous stress ($\mu \frac{d\overline{u}}{dy}$) and Reynolds shear stress ($-\rho \overline{u'v'}$), which arises in single-phase channel flow. The effect of mass loading on the stress balance is shown in figure 2.5. As mass loading increases, the viscous stress stays fairly constant, but the fluid turbulent stress becomes weaker, and is replaced by the particle stress, as previously shown by Richter and Sullivan [169] and Costa et al. [45]. At large mass loading, the wall-normal flux of streamwise momentum is dominated by the particle stress (i.e. the momentum of the particles), which is also an indication of the diminished intensity of turbulence at large mass loading. The case of $\phi_0 = 1$ is of particular interest (figure 2.5(c)). For $y^+ > 100$, the particle and turbulent stresses are nearly equal, which would be expected if particles perfectly follow fluid streamlines.

This interpretation is confirmed by examining the dependence of the stress balance as St^+ is varied with $\phi_0 = 1$ in figure 2.6. For the smallest particle inertia ($St^+ = 1$), the turbulent and particle stresses follow each other closely, except for a small deviation, roughly in the buffer layer. As St^+ increases, this region expands, especially into the viscous sublayer. As a result, the turbulent
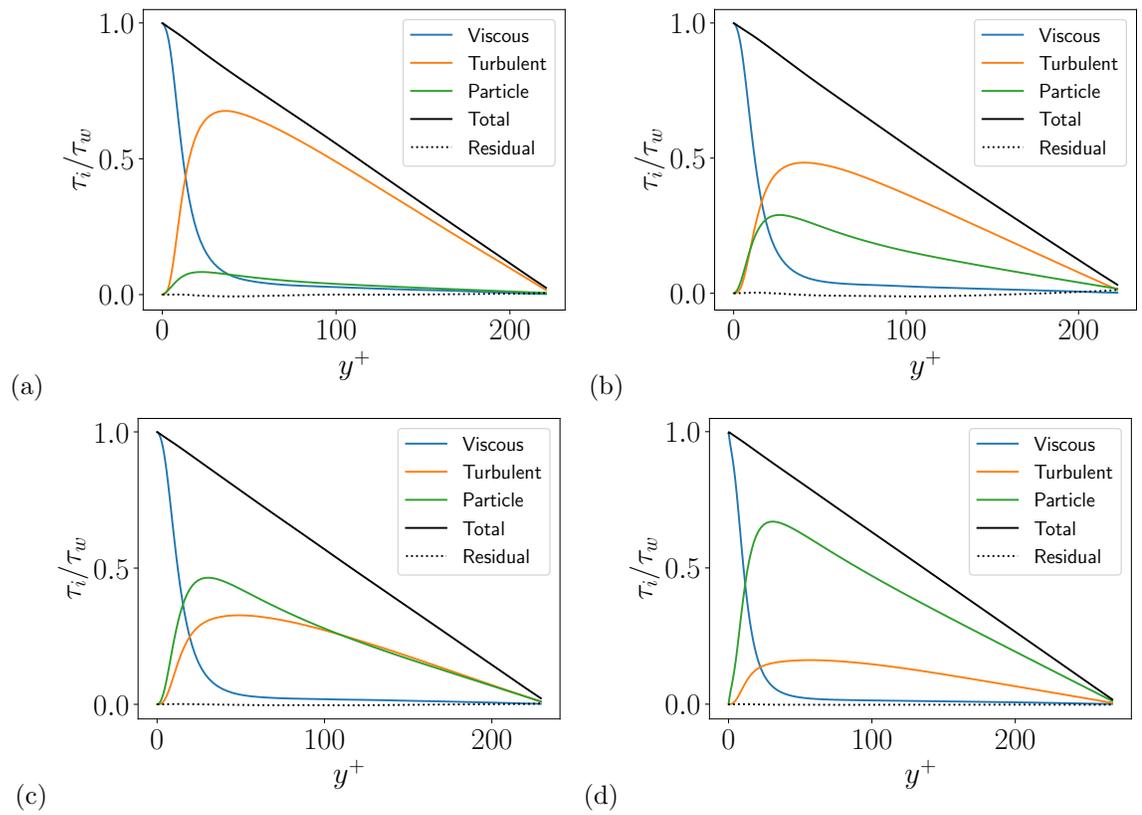
FIGURE 2.5: Streamwise momentum budget as ϕ_0 is varied. (a) $\phi_0 = 0.1$, (b) $\phi_0 = 0.4$, (c) $\phi_0 = 1.0$, and (d) $\phi_0 = 3.0$. The residual (sum) of the terms is plotted with a dotted line.
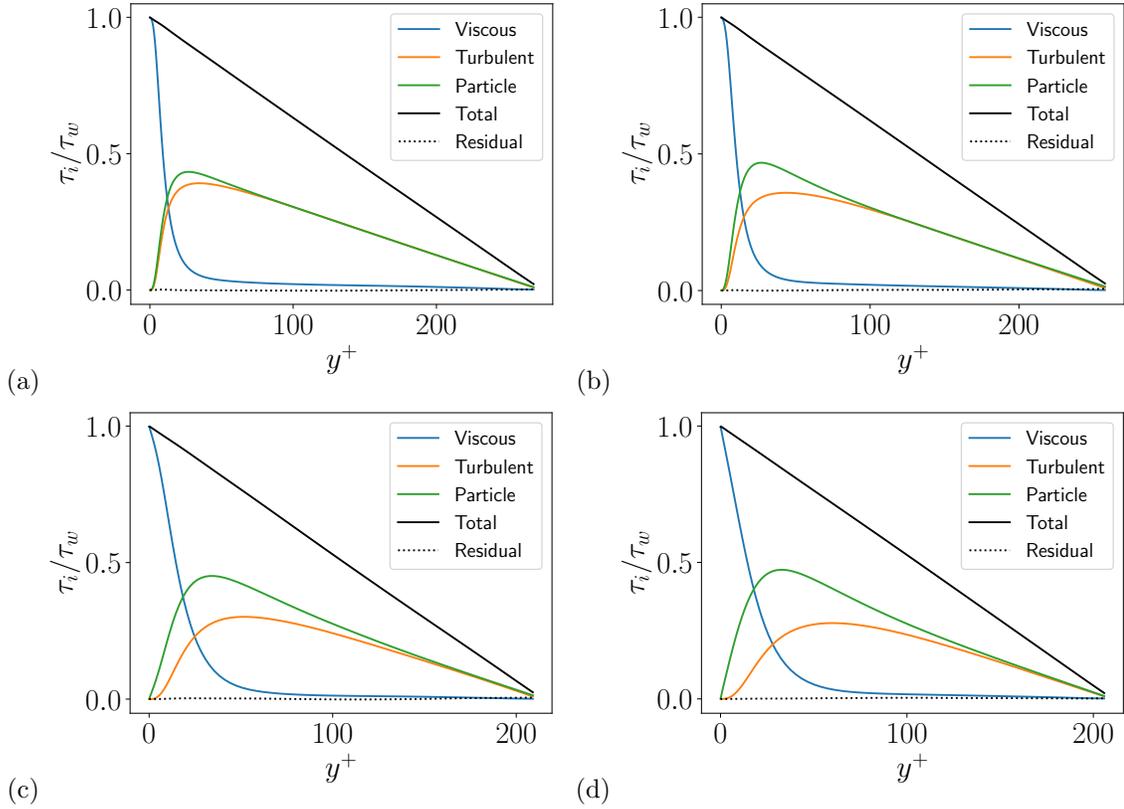
FIGURE 2.6: Streamwise momentum budget as St^+ is varied. (a) $St^+ = 1$, (b) $St^+ = 3$, (c) $St^+ = 20$, and (d) $St^+ = 60$. The residual (sum) of the terms is plotted with a dotted line.

fluid stresses are damped next to the wall, and the viscous stresses decay less rapidly. Interestingly, even at very large Stokes number ($St^+ = 60$), there is a significant chunk of the logarithmic layer, $y^+ \gtrsim 150$, where the particle and fluid stresses are equal, suggesting that particles and fluid are transporting momentum in a more unified way far enough away from the wall. The fact that particle stresses exceed fluid turbulent stresses in this near-wall region is likely due to the particles experiencing a different boundary condition than the fluid, i.e. reflection as opposed to no-slip. As a result, streamwise particle velocity fluctuations do not vanish at the wall, and so their contribution to the particle stress can be greater.

2.3.2   Friction Coefficient

A key bulk flow quantity of interest in wall-bounded flows is the friction coefficient. Depending on the normalization used, different conclusions about drag reduction may be reached, so we include two choices here for comparison: C_f is the friction coefficient traditionally defined based on fluid

bulk velocity, and $C_{f,m}$ is the friction coefficient defined based on a mixture bulk velocity.

$$C_f = \frac{\tau_w}{\frac{1}{2}\rho U_b^2} \tag{2.7}$$

$$C_{f,m} = \frac{\tau_w}{\frac{1}{2}\rho(1+\phi_0)U_{b,m}^2} \tag{2.8}$$

where U_b and $U_{b,m}$ denote the bulk velocities of the fluid and mixture, respectively. These are defined below.

$$U_b = \frac{1}{2\delta\rho} \int_0^{2\delta} (\rho\bar{u})dy \tag{2.9}$$

$$U_{b,m} = \frac{1}{2\delta\rho(1+\phi_0)} \int_0^{2\delta} (\rho\bar{u} + m_p\overline{nu_p})dy \tag{2.10}$$

$$U_{b,m} \approx \frac{1}{2\delta(1+\phi_0)} \int_0^{2\delta} (1+\overline{\phi})\bar{u}dy \tag{2.11}$$

The last equation is an easier to calculate approximation, which omits the effects of particle slip velocity and turbulent mass flux, and is the version used here. Fortunately, the bulk velocities U_b and $U_{b,m}$ are not very different (within 4% across all cases), although the mixture bulk velocity is always smaller than the fluid bulk velocity. This is expected because the particles tend to accumulate in low turbulence regions due to turbophoresis, which also happen to be regions of slower flow near the wall. However, even in the flows with strong turbophoresis, the regions of particle accumulation are very thin and do not comprise much of the overall mass flux.

The friction coefficient as a function of ϕ_0 and St^+ is shown in figure 2.7. C_f shows a minor decrease for $\phi_0 < 1$ before increasing again. This is expected because the centerline velocity in plus units (u_{CL}^+) decreases with mass loading (see figure 2.1). This non-monotonicity was also observed by Capecelatro et al. [35]. in the context of riser flows (gravity acts opposite the flow direction). However, their simulations at high mass loading are dominated by turbulence generated by gravity acting on particle clusters, a mechanism absent in the present work. In simulations without gravity, both drag reduction (such as in Zhao et al. [218]) and drag increase (as in Battista et al. [24]) have been observed with increases in mass loading. In recent particle-resolved simulations at mass loading up to 0.5, Costa et al. [45] observed a small but monotonic increase in skin friction with mass loading, albeit at a slightly lower Reynolds number (Re_τ) than studied here.

However, if the mixture mass flux is accounted for, $C_{f,m}$ is observed to decrease monotonically with mass loading. This decreasing trend is partially captured by the nominal mass loading, but the differences for $\phi_0 > 1$ are due to increasing skin friction. As St^+ is increased at constant mass loading, the friction coefficient decreases most sharply around an intermediate $St^+ \approx 8$, with more gradual decrease at higher and lower St^+. This critical value of St^+ corresponds to $St_\eta^{\log} = 0.54$ (see table 2.1). This is fairly close to the value $St_\eta = 1$, at which preferential concentration peaks
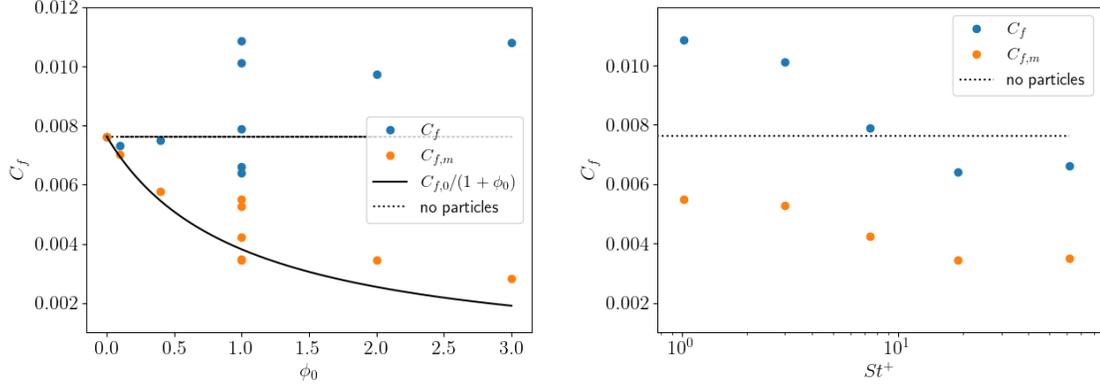
FIGURE 2.7: Drag coefficient for different particle laden flow cases. (**Left**): Effect of Mass loading across all cases. The vertical spread in the data at the same ϕ_0 is due to $St^+ =$ variation. (**Right**): Effect of Stokes number at $\phi_0 = 1$

in homogeneous isotropic turbulence, suggesting that the changes in skin friction may be affected by the degree of particle clustering. These same trends can also be observed by comparing u_{CL}^+ among the cases. In a study of downward particle-laden pipe flow, Vreman [203] found that the friction coefficient normalized based on mixture quantities ($C_{f,m}$ in the present work) decreased for small mass loadings and then increased. This is in contrast to the monotonic decrease observed here, which is mostly likely due to the absence of gravity on particles, which must be balanced by an increase in wall stress in a downward flow.

2.3.3   Fluid turbulent kinetic energy budgets

Next, we examine the effect of particles on the budgets of fluid turbulent kinetic energy (TKE). Like the stress budget, the TKE budget has an additional term due to the particle two-way coupling force. The TKE budget after using homogeneity in x, z, and time, is

$$-\overline{u'v'}\frac{d\overline{u}}{dy} - \frac{d}{dy}\left(\frac{1}{2}\overline{u_i'u_i'v'}\right) - \frac{1}{\rho}\frac{d\overline{p'v'}}{dy} + \nu\frac{d^2}{dy^2}\left(\frac{1}{2}\overline{u_i'u_i'}\right) - \nu\overline{\frac{\partial u_i'}{\partial x_j}\frac{\partial u_i'}{\partial x_j}} + \frac{1}{\rho}\overline{F_i'u_i'} = 0 \qquad (2.12)$$

where the terms, left-to-right, represent shear production, turbulent transport, pressure redistribution, viscous diffusion, viscous dissipation, and the correlation between fluctuations in the fluid velocity and the particle feedback force. This last term may be a source or a sink depending on the relative motion and forces on particles, as discussed by Capecelatro et al. [35]. In compact notation used in the following figures, this equation can be written as

$$\mathcal{P} + \mathcal{T} + \Pi + \mathcal{D} - \epsilon + \mathcal{TWC} = 0 \qquad (2.13)$$

The TKE budgets as mass loading is varied are shown in figure 2.8. All terms are normalized by the viscous scale u_τ^4/ν. Relative to this normalization, most budget terms are attenuated as mass loading increases, except for the two-way coupling term $\mathcal{TWC}$, which becomes more prominent. This term has a similar shape to the shear production term $\mathcal{P}$, and there appears to be a tradeoff between these two mechanisms of production as mass loading increases, similar to the tradeoff between turbulent and particle stress shown in figure 2.5. Among the attenuated terms is the attenuation of pressure redistribution Π, which is responsible for the increasing anisotropy of the turbulence, as described by Vreman [203] and Li et al. [117]. We also note that particles change the production-dissipation balance in the logarithmic layer. As mass loading increases, this shifts to a balance among shear production and the two-way coupling source on the positive side, and viscous dissipation on the negative side. Finally, we note that the residual is non-zero, and can be comparable in magnitude to the pressure redistribution and turbulent transport terms. This is due to the resolution of the simulations, and is shown in Appendix 2.6 to reduce with grid refinement. No term changes substantially, and the interpretation of the above is not affected.

Figure 2.9 shows how the TKE budget changes as St^+ is varied at $\phi_0 = 1$. For small St^+, the shear production and two-way coupling source follow each other more closely, while other terms are not significantly affected. As St^+ increases, most terms are slightly attenuated, and the peaks of terms are stretched away from the wall, which further indicates a thickening of the viscous sublayer. For $St^+ = 60$, the two-way coupling term becomes slightly negative in the log layer, indicating that fluid TKE is dissipated via interactions with the particles. Interestingly, the residual shrinks as St^+ increases, suggesting that the grid used is of sufficient resolution when the viscous sublayer is slightly thickened relative to the unladen case.

2.4 Scaling and Transforms

The results in the preceding section indicate that the traditional scaling used for incompressible channel flow do not result in a collapse of the streamwise velocity data, and the momentum and TKE budgets suggest that other scalings might be suitable. Here we consider a transform inspired by work on compressible flow transforms, especially Trettel and Larsson [198] and Patel et al. [158].

We begin this derivation by considering the equations for the momentum of the fluid and particle phases, in Eulerian form:

$$\frac{\partial (\rho \boldsymbol{u})}{\partial t} + \nabla \cdot (\rho \boldsymbol{u} \otimes \boldsymbol{u}) = -\nabla p + \mu \nabla^2 \boldsymbol{u} + \boldsymbol{F}, \tag{2.14}$$

$$\frac{\partial (m_p n \boldsymbol{v}_p)}{\partial t} + \nabla \cdot (m_p n \boldsymbol{v}_p \otimes \boldsymbol{v}_p) = -\boldsymbol{F}, \tag{2.15}$$

where n is the particle number density, defined on an Eulerian grid. Here, the particle velocity $\boldsymbol{v}_p$ is intended as the spatially correlated component of the particle velocity [67], which is obtained by
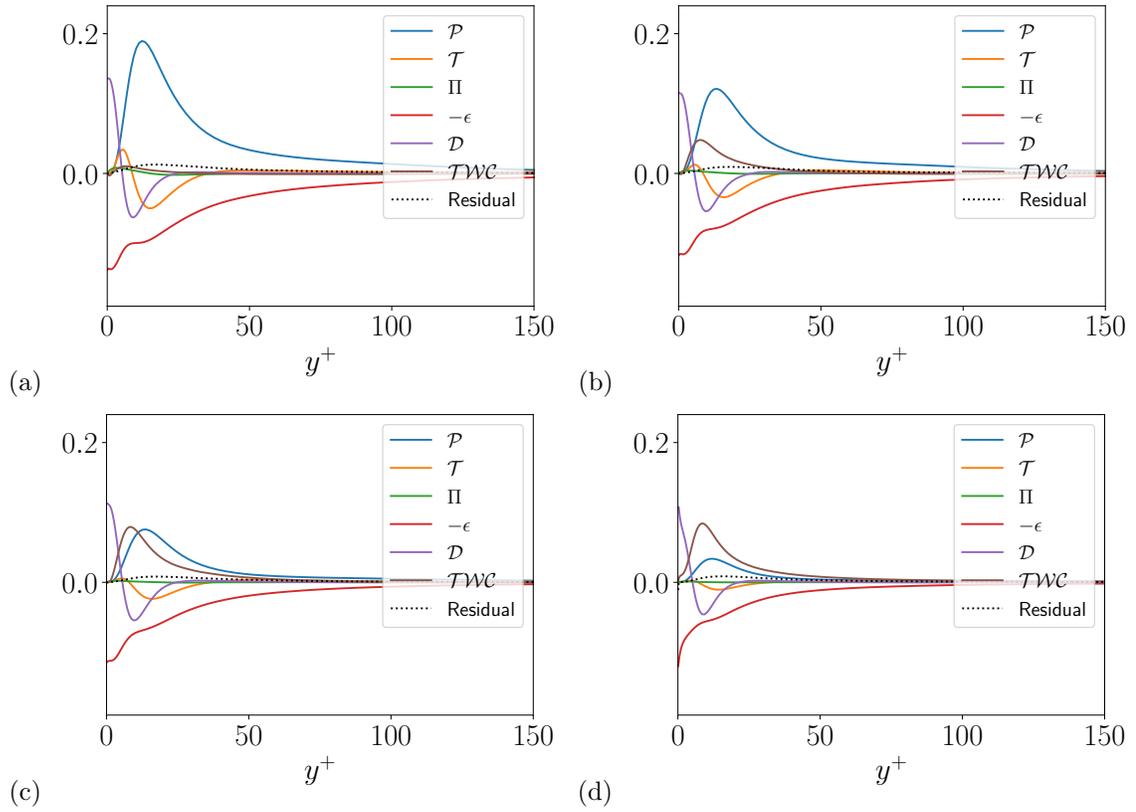
FIGURE 2.8: Turbulent kinetic energy budget as ϕ_0 is varied, normalized by viscous scales. Each line corresponds to a term in equation 2.13. The residual (sum) of the terms is plotted with a dotted line. (a) $\phi_0 = 0.1$, (b) $\phi_0 = 0.4$, (c) $\phi_0 = 1.0$, and (d) $\phi_0 = 3.0$.
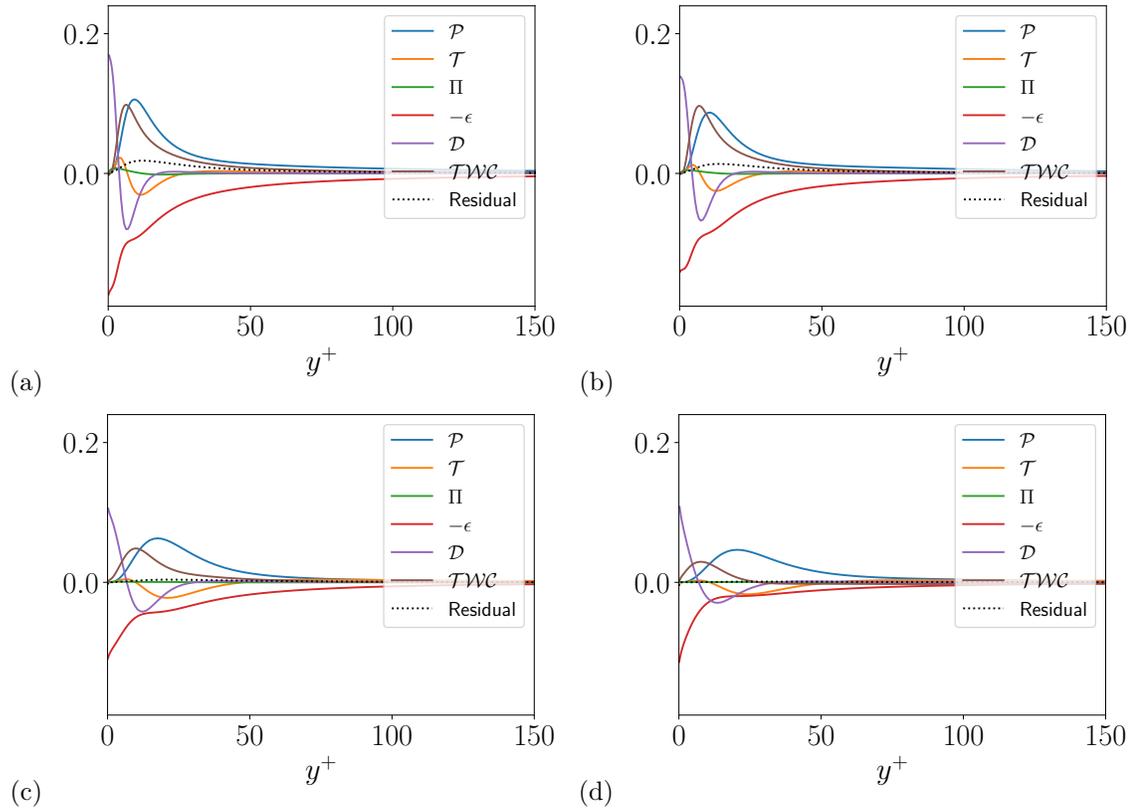
FIGURE 2.9: Turbulent kinetic energy budget as St^+ is varied, normalized by viscous scales. Each line corresponds to a term in equation 2.13. The residual (sum) of the terms is plotted with a dotted line. (a) $St^+ = 1$, (b) $St^+ = 3$, (c) $St^+ = 20$, and (d) $St^+ = 60$.

averaging the particle velocities over their velocity distribution in a computational cell. We sum these equations, and average in time and in the homogeneous spatial directions. The resulting equation for streamwise mixture momentum is

$$(1 - \frac{y}{\delta}) = \frac{1}{\tau_w} \left(\mu \frac{d\overline{u}}{dy} - \left[\rho \overline{u'v'} + m_p \overline{n u_p v_p} \right] \right),$$ (2.16)

Note that this equation is the same as equation 2.6, except that the particle stress has been written as a particle momentum flux. This "particle Reynolds stress" term includes both the mean and fluctuating number density n. The particle momentum terms can be expanded and somewhat simplified by considering the Eulerian continuity equation for particle mass:

$$\frac{\partial (m_p n)}{\partial t} + \nabla \cdot (m_p n \boldsymbol{v}_p) = 0.$$ (2.17)

Averaging in time and homogeneous spatial directions gives

$$\frac{d}{dy} \left(\bar{n} \bar{v}_p + \overline{n' v_p'} \right) = 0.$$ (2.18)

Because of the no penetration boundary condition, there can be no net mass flux of particles in the y direction once steady state is reached. Therefore the mean and fluctuating particle velocities are related according to

$$\bar{n} \bar{v}_p + \overline{n' v_p'} = 0$$ (2.19)

Using this result, the mixture x-momentum equation can be written

$$(1 - \frac{y}{\delta}) = \frac{1}{\tau_w} \left(\mu \frac{d\overline{u}}{dy} - \left[\rho \overline{u'v'} + m_p \left(\bar{n} \overline{u_p' v_p'} + \bar{v}_p \overline{n' u_p'} + \overline{n' u_p' v_p'} \right) \right] \right),$$ (2.20)

Now, we make use of the fact that the particle velocity and fluid velocities are related by $u_i{}^{(p)} = u_i + \mathcal{O}(St^+)$. Taking the zero-inertia limit, which should be applicable to small Stokes numbers, the resulting approximate equation is

$$(1 - \frac{y}{\delta}) \approx \frac{1}{\tau_w} \left(\mu \frac{d\overline{u}}{dy} - \left[\rho \overline{u'v'} + m_p \left(\bar{n} \overline{u'v'} + \bar{v} \overline{n' u'} + \overline{n' u'v'} \right) \right] \right),$$ (2.21)

We further note that $\overline{v} = 0$ due to fluid incompressibility, and assume that the triple correlation $\overline{n' u' v'}$ is small, to obtain

$$(1 - \frac{y}{\delta}) \approx \frac{1}{\tau_w} \left(\mu \frac{d\overline{u}}{dy} - \bar{\rho}_m \overline{u'v'} \right),$$ (2.22)

where ρ_m is the mixture density $\rho_m = \rho + m_p n$. Unlike the fluid density, the mixture density varies in time and space, due to preferential concentration and turbophoresis. Equation 2.22 suggests that the the momentum transfer in particle-laden wall-bounded flow will be similar to a variable
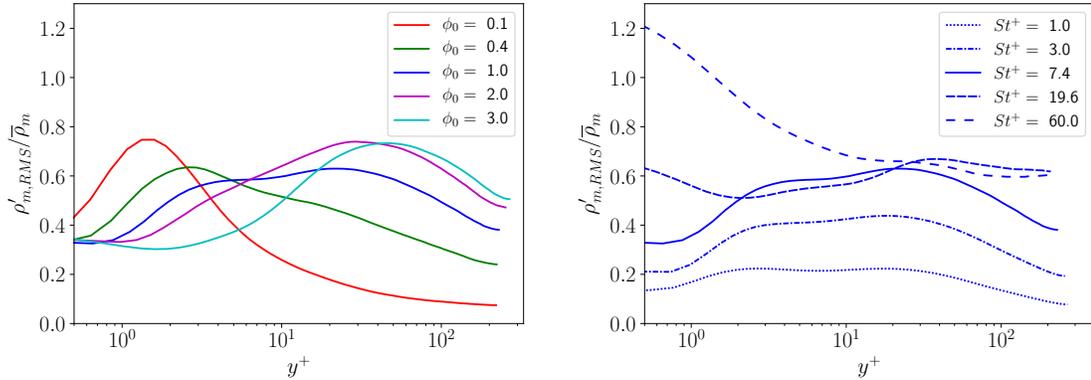
FIGURE 2.10: Strength of mixture density fluctuations relative to the mean value, spanning variation in ϕ_0 (left) and St^+ (right).

density flow, where the viscosity is the fluid viscosity and the density is the mixture density. This suggests that transform approaches previously applied to variable density and compressible flows, such as Van Driest [201], and more recently Trettel and Larsson [198] and Patel et al. [158] may be applicable to particle-laden flows, with these substitutions.

Existing transform approaches have only been found to be effective when compressibility effects are sufficiently small. In particle-laden flows, even though the fluid may be incompressible, the mixture has variable and fluctuating density. Furthermore, the instantaneous particle velocity field may be multi-valued (i.e. containing caustics), which will lead to qualitative differences in the Reynolds stresses which transport velocity toward the wall. Caustics are known to become more prevalent as particle inertia increases [209]. For all of these reasons, it seems likely that St^+ and ϕ_0 may play a role in determining the transform applicability, in much the same way that a turbulent Mach number does for compressible flows. Figure 2.10 shows the strength of number density fluctuations as a function of mass loading and Stokes number. This shows that the mixture density fluctuations in the dataset considered here are substantial, but generally smaller than the mean mixture density, except for very large values of St^+. The fluctuations are generally observed to increase with increasing Stokes number. As mass loading is increased at moderate Stokes number, there is a peak in fluctuations that moves outward toward the centerline, and the intensity at the centerline increases, but the overall fluctuation strength is less affected by ϕ_0 than by St^+.

2.4.1   Assessment of transform approaches

The preceding derivation suggests that the Van Driest or Trettel-Larsson transforms may be applicable to particle-laden flows under certain conditions. For this purpose, we define a mixture density at the wall $\rho_{m,w}$, a mixture friction velocity $u_{\tau,m} = \sqrt{\tau_w/\rho_{m,w}}$, and a mixture velocity normalization

$u_m^+ = u/u_{\tau,m}$. With these substitutions, the "mixture Van Driest" transform is defined as

$$y_{VD,m}^+ = y^+ \tag{2.23}$$

$$u_{VD,m}^+ = \int_0^{u_m^+} \left(\frac{\bar{\rho}_m}{\rho_{m,w}}\right)^{\frac{1}{2}} du_m^+. \tag{2.24}$$

where mixture density and mixture friction velocity have been substituted for their single-phase counterparts. Similarly, we define the "mixture Trettel-Larsson transform" as

$$Y_m^* = \frac{\bar{\rho}_m (\tau_w/\bar{\rho}_m)^{\frac{1}{2}} y}{\bar{\mu}} \tag{2.25}$$

$$U_m^* = \int_0^{u_m^+} \left(\frac{\bar{\rho}_m}{\rho_{m,w}}\right)^{\frac{1}{2}} \left[1 + \frac{1}{2}\frac{1}{\bar{\rho}_m}\frac{d\bar{\rho}_m}{dy}y - \frac{1}{\bar{\mu}}\frac{d\bar{\mu}}{dy}y\right] du_m^+. \tag{2.26}$$

The result of applying these transforms is shown in figure 2.11. An encouraging sign is that across most of the cases, the slope in the log layer of the transformed velocity profiles now closely matches the slope in the incompressible profile, but an offset is introduced. This can be seen in a vertical shift (in U_m^*) in the log layer as well as a horizontal shift (in Y_m^*) in the start of the log layer. This is reminiscent of a roughness effect, which manifests as an offset due to the addition of a length scale in the near-wall region [193]. It is also worth noting the amount of offset increases with both ϕ_0 and St^+. The fact that the transform collapse becomes worse as ϕ_0 or St^+ increases is expected, as the strength of mixture density fluctuations (figure 2.10) increases with both of these quantities.

We quantify the above observations with three characteristics of the streamwise velocity profile:

1. The viscous sublayer slope S, shown in figure 2.12.

2. The log layer slope $1/\kappa$, shown in figure 2.13

3. The log layer offset C, shown in figure 2.14

The logarithmic layer parameters are calculated by finding the κ and C which minimize the following functional for candidate velocity and coordinate $u^* = u^*(y^*)$

$$\int_{y^*=60}^{y^*=\delta^*} \left(u^*(y^*) - \left(\frac{1}{\kappa}\log(y^*) + C\right)\right)^2 dy^* \tag{2.27}$$

where δ^* signifies integration to the centerline. To collapse data to incompressible channel flow, the desired behavior is $S = 1$, $1/\kappa \approx 2.5$, and $C \approx 5$ [193].

In terms of the viscous sublayer slope, shown in figure 2.12, both the traditional viscous scaling and the mixture Trettel-Larsson transform are adequate. The mixture Van Driest produces the incorrect slope, which has long been observed in the context of supersonic boundary layers with cooled walls [198]. The desired log layer slope, shown in figure 2.13, is better captured by both
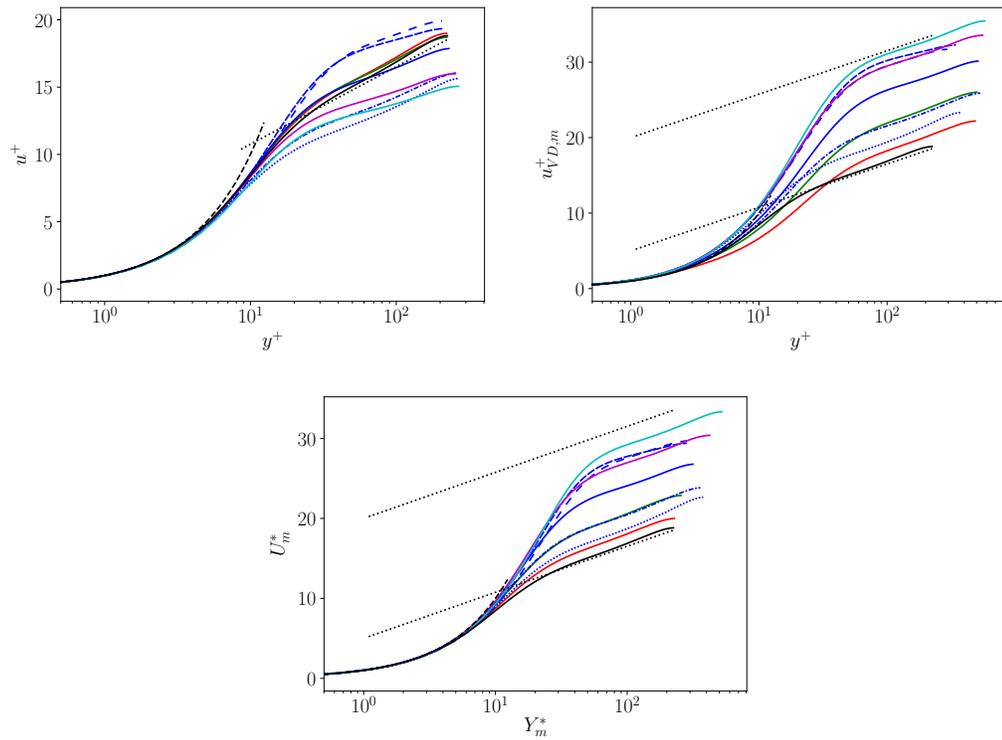
FIGURE 2.11: Mean fluid velocities obtained using the (a) traditional viscous scaling, (b) the "mixture Van Driest " transform, and (c) the "mixture Trettel-Larsson" transform, for simulation cases spanning variation in ϕ_0 (indicated by line color) and St^+ (indicated by line style). Line styles and colors are the same as in figure 2.10. The near-wall viscous scaling is plotted in all sub-figures, and two logarithmic profiles with different offsets are plotted for reference in (b) and (c).
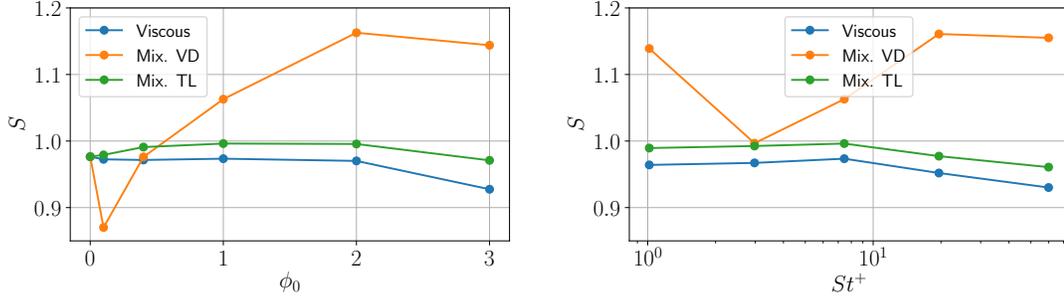
FIGURE 2.12: Viscous sublayer slope for various transform approaches, plotted against the variation in mass loading (left) and Stokes number (right).
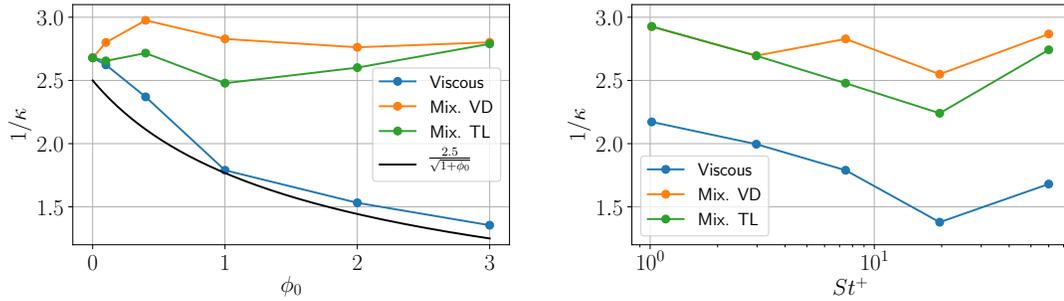
FIGURE 2.13: Logarithmic layer slope for various transform approaches, plotted against the variation in mass loading (left) and Stokes number (right).

mixture transformations. There does seem to be some non-monotonic dependence on St^+ for the log layer slope, which may be a consequence of preferential concentration, since particles are not following fluid streamlines perfectly, as assumed in the transform derivation. The observed variability of the log layer slope with ϕ_0 using the traditional scaling can be nearly predicted by the formula $2.5/\sqrt{1 + \phi_0}$, which can be justified by assuming that the mixture density in equation 2.22 is uniform, and applying the derivation in Tennekes and Lumley [193] to obtain the log layer profile. Finally, we note that the offset, shown in figure 2.14, increases regardless of transform used, but that it increases more for the mixture transformations. This offset increases as a function of mass loading and Stokes number. Both of the quantities suggest a saturation in the amount of the offset needed, but this is not certain.

2.4.2   Roughness-like offset

Noting the mathematical similarity to a roughness offset, an attempt to further collapse the data can be made by introducing an offset into the mixture Trettel-Larsson transform. First, we define
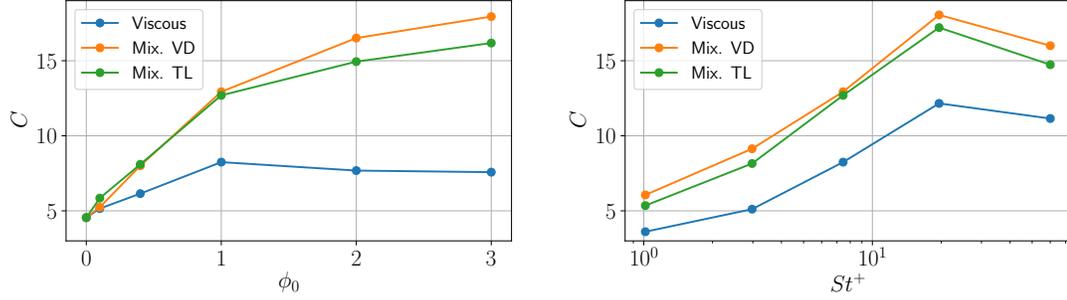
FIGURE 2.14: Logarithmic layer offset for various transform approaches, plotted against the variation in mass loading (left) and Stokes number (right).

an offset d and an offset coordinate s

$$s = y - d. \tag{2.28}$$

This leads to definitions for an offset wall stress τ_d and modified viscous scales.

$$\tau_d = \tau_w \left(1 - \frac{d}{h}\right) \tag{2.29}$$

$$u_{m,d} = \sqrt{\frac{\tau_d}{\rho_{m,d}}} = u_{m,\tau} \sqrt{\left(1 - \frac{d}{h}\right)} \tag{2.30}$$

$$u_{m,d}^+ = \frac{u}{u_{m,d}} \tag{2.31}$$

which are used to define a mixture-based, offset Trettel-Larsson transform.

$$Y_{m,offset}^* = \frac{\bar{\rho}_m (\tau_d/\bar{\rho}_m)^{\frac{1}{2}} s}{\bar{\mu}} \tag{2.32}$$

$$U_{m,offset}^* = \int_0^{u_{m,d}^+} \left(\frac{\bar{\rho}_m}{\bar{\rho}_{m,d}}\right)^{\frac{1}{2}} \left[1 + \frac{1}{2}\frac{1}{\bar{\rho}_m}\frac{d\bar{\rho}_m}{ds}s - \frac{1}{\bar{\mu}}\frac{d\bar{\mu}}{ds}s\right] du_{m,d}^+ \tag{2.33}$$

Given the above formulation, offsets can be found empirically such that collapse in the log layer is obtained. This is done by finding the offset which minimizes the error with respect to the incompressible profile in the range $30 < y^+ < 150$. Following this procedure, the resulting offsets are plotted versus the parameter space in figure 2.15, using two different normalizations. These show that the offset heights are the same order of magnitude as the mean inter-particle distance based on the mean number density at the wall, as well as the viscous length scale.

The viscous scaling suggests that the mechanism for the offset may be the thickening of the viscous sublayer. This is an appealing interpretation because of the observed attenuation in fluid Reynolds shear stress, but it is complicated by the fact that streamwise velocity fluctuations remain strong (figure 2.3). With the viscous scaling, the offset required for collapse increases with both ϕ_0
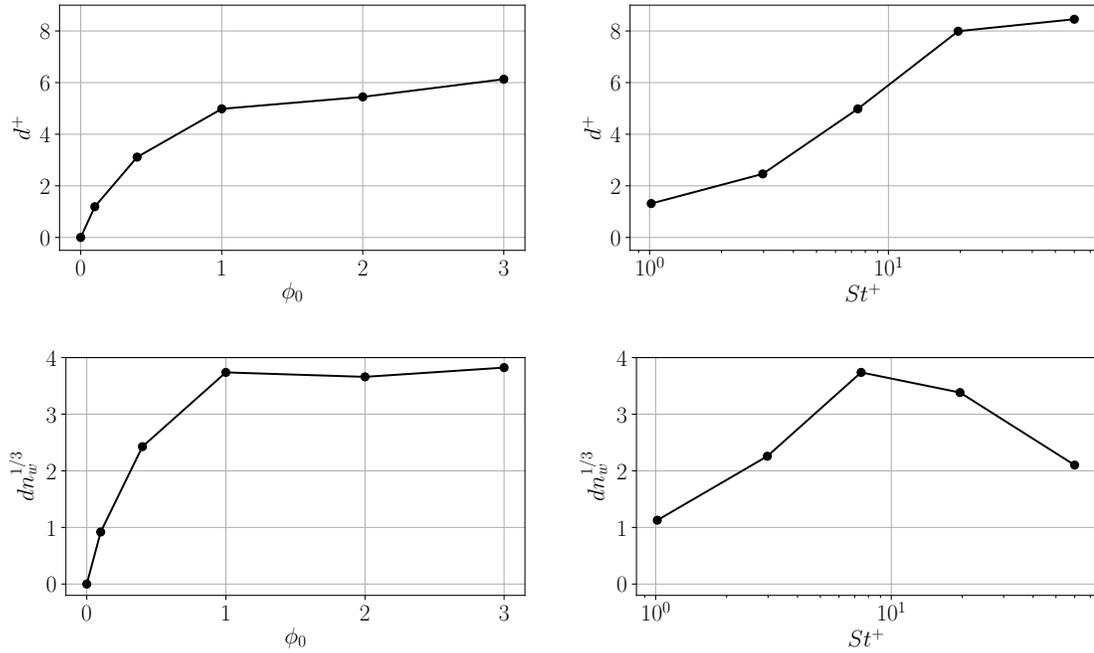
FIGURE 2.15: Empirically determined offsets plotted against mass loading (left) and Stokes number (right). The top row uses the viscous scales for normalization. The bottom row uses a normalization of mean inter-particle distance based on the mean near-wall number density.

and St^+, seeming to saturate at large values.

The scaling based on mean inter-particle distance suggests a different interpretation. It is expected that particle clusters disrupt eddies similar to their size: smaller eddies will pass between particles, while larger eddies will move whole cluster together. As a result, the presence of particles may disrupt the structure of turbulence in a thin layer near the wall, where eddy size scales with wall-distance. This interpretation is appealing because the offset saturates at large mass loading, and it is observed that particle mass loading at the wall does the same in figure 2.2. With Stokes number, there is a non-monotonic trend, with the largest offset required at the intermediate value $St^+ = 7$, where preferential concentration is the greatest.

If the behavior observed is akin to a roughness offset in rough-wall flows, then it implies that there is a near-wall region where the turbulence behavior is altered by the presence of particles, the thickness of which is an additional length scale imposed on the flow. Costa et al. [43] described such a layer for a dense mixture with neutrally buoyant particles, in which particle slip velocity next to the wall was large. This led to an offset thickness proportional to the particle diameter. The physical explanation must be different for the present study because the mixtures are dilute, slip velocities are observed to be small, and the offsets needed are in the range of $10 < d/D_p < 70$.
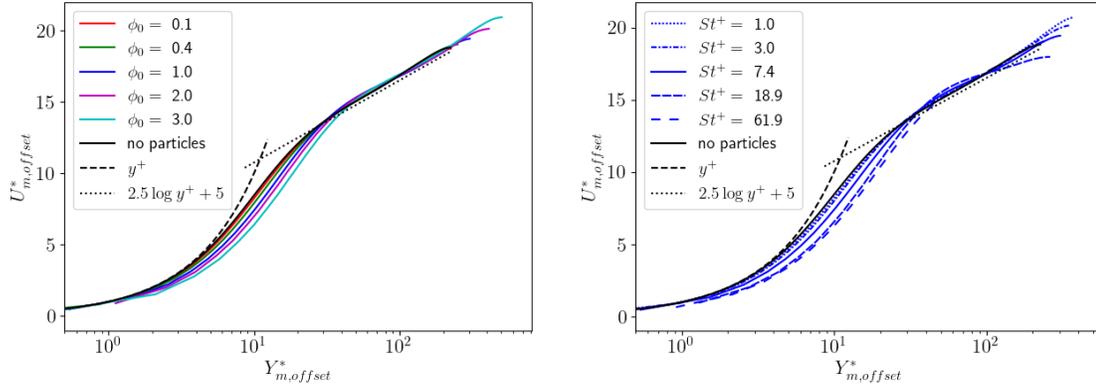
FIGURE 2.16: Velocity profiles obtained by the mixture Trettel-Larsson transform with empirical offset, plotted across variation in mass loading (left) and Stokes number (right)

Using the empirically determined offsets, reasonable collapse for velocity profiles can be obtained in the log layer, as shown in figure 2.16. The collapse is better for lower St^+, and quite good for the variation in ϕ_0. Finally, as the offset height is increased, the profiles align more poorly in the viscous sublayer and buffer layer, which is expected.

2.5 Conclusion

In this work we have described turbulence modulation in channel flow by inertial particles, and made progress toward a velocity transform for particle-laden wall-bounded flows by borrowing ideas from the study of compressible turbulence. The effect of mass loading and Stokes number on channel flow has been shown, in agreement with other studies. The addition of particles damps wall-normal fluid velocity fluctuations and slightly enhances streamwise ones. This leads to an attenuation of Reynolds shear stress, and a greater role played by the particle stress as mass loading and Stokes number are increased. Momentum budgets reveal that for small Stokes number, the particle stress carries the proportionate part of the fluid stress, in accordance with the mass loading. As Stokes number increases, particle stresses exceed fluid stresses near the wall, and the region over which particle stresses are larger expands as Stokes number increases. Similar effects are seen in the shear production term in the turbulent kinetic energy (TKE) budget, due to its dependence on the Reynolds stress. The presence of particles damps most TKE budget terms, but there is a new term related to the particle phases, which behaves similar to fluid shear production, except at large enough Stokes number, when it takes on a dissipative character. In the log layer and centerline, traditional production-dissipation balance is augmented with the particle exchange term, which becomes more significant as mass loading is increased.

In terms of transform approaches, we have demonstrated how an improved collapse of particle

	Re_τ	N_p	ρ_p/ρ	ϕ_0	ϕ_{visc}	ϕ_{buf}	ϕ_{log}	St^+	St_η^{log}
Flow 0	229	0	N/A	0	0	0	0	N/A	N/A
Flow 0 refined	228	0	N/A	0	0	0	0	N/A	N/A
Flow 3	234	1.43×10^8	7.42×10^3	1.0	3.54	1.41	0.88	7.45	0.54
Flow 3 refined	236	1.43×10^8	7.42×10^3	1.0	3.76	1.47	0.86	7.60	0.59

TABLE 2.2: Parameters of the channel flow data with grid refinement: the friction Reynolds number Re_τ, the number of particles N_p, the particle to fluid density ratio ρ_p/ρ, the overall mass loading ϕ_0, the average mass loading in the viscous sublayer ϕ_{visc}, the average mass loading in the buffer layer ϕ_{buf}, the average mass loading in the logarithmic layer ϕ_{log}, the friction Stokes number St^+, and the Kolmogorov Stokes number based on the dissipation rate averaged over the logarithmic layer St_η^{log}.

laden flow velocity data in the dilute but strongly coupled regime can be obtained by appealing to ideas from compressible flow transforms and rough wall turbulent flows. Accounting for the particle contribution to the mixture density in a Trettel-Larsson type transform is effective in collapsing the slopes of velocity profiles in log layer. However, this introduces a roughness-like offset that increases with mass loading and Stokes number, but the physical mechanism underlying this roughness type offset remains to be explained. Nevertheless, it was verified that there do exist offset values that produce good collapse of data in the log layer, but they are empirically tuned for this purpose.

In the transformed profiles, the offset appears in the buffer layer and reaches a constant value in the log layer, suggesting that the transform is adequate except in the buffer layer. This difficulty may be due to the strong turbulent fluctuations in the buffer layer, which is home to maximum TKE, TKE production, and TKE turbulent transport [163]. The relative intensity of the turbulence in the buffer layer may lead to particles not following fluid streamlines as closely, and become more independent of fluid eddies, which undermines the reasoning used in developing the transform. In this chapter, we have focused on the statistics of the fluid phase, but to better understand the partial success of the transform, as well as gain a more holistic understanding of the dynamics of the mixture, it is useful to study the particle dynamics. These are considered in the next chapter, which examines clustering, vorticity, and swirling motions in the particle phase.

The focus of this work has been isothermal particle-laden flow, in which compressibility-like effects occur due to particle distribution. The framework laid out here can naturally incorporate variability in fluid density and viscosity as well, which will occur in particle laden flows with heat transfer, such as dusty re-entry flows and solar energy applications. Finally, we also note that in principle the variable density transform ideas proposed here can be combined with the scaling from Costa et al. [43], in order to better collapse velocity data from complex applications in which particles modify both the viscosity and density of the mixture.

2.6   Appendix A: Grid convergence

To assess the grid convergence of the results presented here, we performed two additional simulation cases with double the grid points in each direction, for the single-phase case (Flow 0) and the case with $St^+ = 7$, $\phi_0 = 1$ (Flow 3). Key flow parameters are summarized in table 2.2. This shows that with grid refinement, there are small changes in Reynolds number, Stokes number, and distribution of mass loading. Next we present statistics and budgets comparing the grids used in the body of the manuscript to their refined counterparts. There are small quantitative changes, but these do not change any of the interpretation of results presented in the body of this work.

In figure 2.17, we show grid convergence behavior for the single-phase channel flow by reporting various statistics with both the baseline grid used in the body of this work, as well as the refined grid. These show that the streamwise velocity is well-resolved and that the turbulent kinetic energy is slightly under-resolved with the baseline grid. The Reynolds stresses and viscous stresses are well-converged. The TKE budget residual decreases noticeably for the finer grid. Not surprisingly, this is due largely to changes in the viscous dissipation $(-\epsilon)$ and turbulent transport terms $(\mathcal{T})$, which are dominated by fine-scale turbulent fluctuations.

The changes with grid refinement for the $\phi_0 = 1$, $St^+ = 7$ case are shown in figure 2.18. These are quantitatively of a similar level to the changes to the unladen flow. Both the streamwise velocity and TKE were slightly over-estimated by the baseline grid. The balances of stresses shifts very slightly from the turbulent fluid stress to the particle phase with grid refinement. This results in a quantitative change to the shear production $\mathcal{P}$ terms in the TKE budget, most likely due to the dependence of this terms on the Reynolds shear stress. The two-way coupling source $\mathcal{TWC}$ is similarly modified. Like the unladen flow, the residual improves with grid refinement, although in this case it is due to the combination of better-resolved viscous dissipation and changes to shear production.
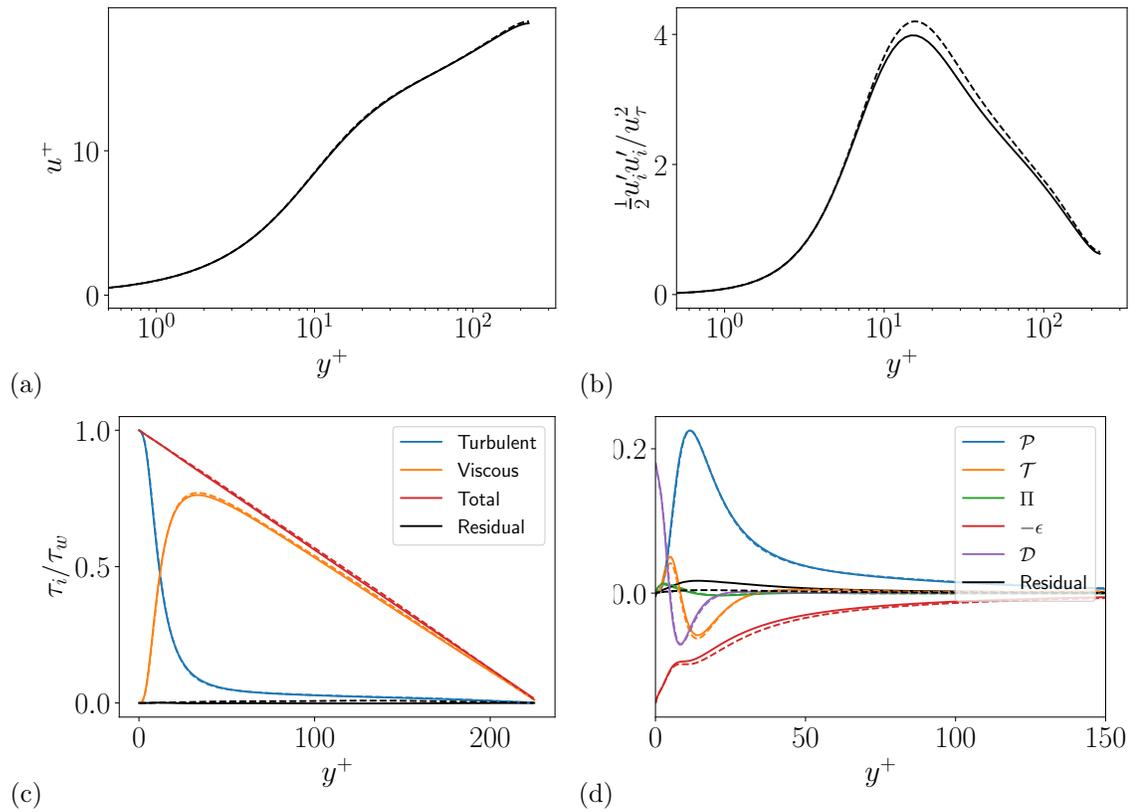
FIGURE 2.17: Effect of grid refinement on various statistics for the single-phase channel flow. Solid lines show the baseline grid $[280, 140, 140]$ and dashed lines show the refined grid $[560, 280, 280]$ for (a) streamwise velocity, (b) turbulent kinetic energy, (c) stress budget, and (d) TKE budget.
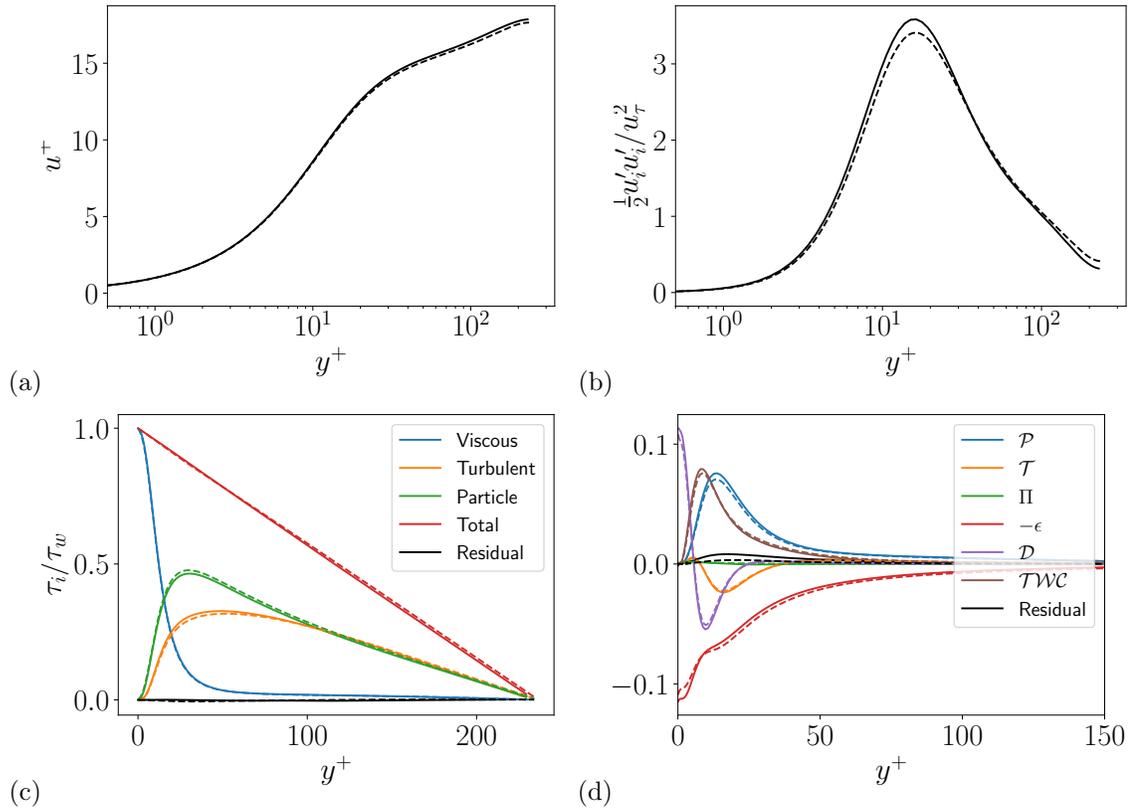
FIGURE 2.18: Effect of grid refinement on various statistics for the case with $\phi_0 = 1$, $St^+ \approx 7$. Solid lines show the baseline grid $[280, 140, 140]$ and dashed lines show the refined grid $[560, 280, 280]$ for (a) streamwise velocity, (b) turbulent kinetic energy, (c) stress budget, and (d) TKE budget.

Chapter 3

Clustering, rotation, and swirl of particle clouds in turbulent channel flow

With the exception of some minor changes, the contents of this chapter have been submitted to the the Journal of Fluid Mechanics [207], and were also published in a preliminary version in the 2022 Proceedings of Center for Turbulence Research Summer Program [205].

3.1 Introduction

Particle-laden turbulent flows over walls occur in many natural and engineering systems. Examples include sediment transport in rivers, dust storms in the atmospheric boundary layer [91], chemical processes in industrial riser reactors [23], and novel radiation absorbers in concentrating solar power plants [84]. In all of these examples, particles are transported by the flow and exert drag forces back on the flow. Experiments by [113] showed that at sufficiently high mass loading, the particle phase can significantly attenuate the fluid turbulence in channel flow. Other aspects of turbulence modulation, including enhanced dissipation by particles, inter-phase energy transfer, and vortex shedding by particles have been studied as a function of particle inertia, mass loading, and other flow parameters [13, 31].

Because of their finite inertia, particles preferentially concentrate in regions of low vorticity and high strain, leading to complex dynamics of cluster formation [129, 185]. Despite this simple description, clustering is known to be weak at very small and very large Stokes numbers and to peak at intermediate Stokes numbers. Recent work by Esmaily-Moghadam and Mani [57] and [56]

has succeeded in predicting this trend analytically by analysing the Lyapunov exponents of small-scale particle clouds. Sweeps and ejections in wall turbulence also play an important role in cluster dynamics and near-wall accumulation in particle-laden channel flow [123]. As a result, these clusters are constantly formed and annihilated as well as subjected to strong rotational and swirling motions near the wall. Clustering also affects the optical depth in gas flows laden with particles, which is particularly important in the design of particle-based solar receivers [64] and in radiation absorption in clouds [126]. The complex physics of cluster formation, combined with the strong turbulence modulation which occurs at high mass loading, makes these flows difficult to predict and model [13]. In addition, the large optical depth makes them difficult to study experimentally, making simulations an essential tool to improve understanding.

In the present work, we study the clustering, rotation, and swirl of particle clouds in a turbulent channel with dilute volume fraction, but with mass loading $\mathcal{O}(1)$, making the effect on the carrier phase dynamics significant. These flows are impractical to simulate with particle-resolved direct numerical simulations, so we opt for an Euler-Lagrange approach [112], which allows for exploration of a wide range of mass loading and Stokes number at affordable cost. The large number of particles also makes analysis of particle clustering difficult, motivating the use of new methods for doing so.

A useful numerical tool for analysing particle clustering in Euler-Lagrange simulations is tessellation. In brief, a tessellation decomposes a domain into geometric shapes. This is useful for analysing particle clustering by defining a tessellation based on the particle locations. Monchaux et al. [138] used the Voronoi tessellation to characterise the effect of Stokes number on particle clustering. This technique has been widely applied in homogeneous isotropic turbulence, and to a lesser extent in wall-bounded flows. For example, [145] confirmed the non-monotonic behaviour of clustering with respect to Stokes number in channel flow using Voronoi tessellation in wall-parallel slices of the domain. To gain further insight into the dynamics of particle clouds, we have recently developed tessellation-based techniques to compute velocity gradients from dispersed particle data. Oujia et al. [151] proposed a method to quantify the formation of clusters and voids by computing the time change of the tessellation cell volume. A method for computing the curl of the particle velocity was also introduced in [152, 153], in order to quantify the rotation of particle clouds.

Knowing the curl, the helicity of the particle motion can also be determined, which yields insight into the processes that determine the flow topology [137]. Swirling flows are characterised by strong helicity, and in the turbulence community, helicity has been used to characterise three-dimensional swirling coherent structures, which correspond to flow regions of maximum helicity [58]. Similarly, the helicity of the particle velocity can be used to determine whether particle motions are two-dimensional (zero helicity) or three-dimensional.

The goal of the present work is to bring detailed insights into the clustering and motion of heavy, inertial particles in wall turbulence, in the regime of moderate mass loading. This is accomplished by

applying tessellation-based analysis to a set of numerical data of high-fidelity, four-way coupled Euler-Lagrange simulations. In the analysis, we address both Eulerian (fluid) and Lagrangian (particle) data. To provide context for the distributions of divergence, curl, and helicity, we first present Eulerian statistics for the fluid phase, including mean velocity profile, mass loading, turbulent kinetic energy (TKE), and Reynolds shear stress against the wall unit. Unlike in HIT, the cluster dynamics in channel flow naturally have a strong dependence on their wall-normal distance. In order to quantify this dependence, we compute the PDFs of divergence, curl, and helicity in the viscous sublayer, buffer layer, and logarithmic layers for both fluid and particle phases. To the best of our knowledge, such an analysis has not been previously performed.

The PDFs of curl and helicity are shown for the fluid phase to illustrate the effect of the particle coupling on rotational and swirling motions of the fluid, and to provide a baseline against which the Lagrangian particle data is compared. The PDFs for the Lagrangian data are reported to characterise cluster/void formation, rotation, and swirling motions of the particle clouds. The effects of particle mass loading and Stokes number on the PDFs of each layer are characterised with the aid of flow visualisations obtained by colouring particles with divergence, vorticity magnitude, or helicity. In the log layer, in which turbulence becomes more homogeneous and isotropic than the other layers, differences and similarities with the previous results of HIT are discussed with comparisons to the work of Matsuda et al. [127] and Oujia et al. [151]. Moreover, the particle-based Fourier transform is used to obtain scale information on the particle number density distribution, and to understand what scales of clustering contribute to the trends observed in the PDFs. A preliminary version of this study was reported in West et al. [205].

The remainder of this manuscript is organised as follows. In § 3.2, the numerical methods used in this work are described, including the channel flow simulation database (§ 3.2.1), the tessellation approach for obtaining divergence, curl, and helicity (§ 3.2.2), and the method for calculating Fourier spectra of particle number density fluctuations (§ 3.2.3). Section 3.3 presents Eulerian statistics for the fluid phase (§ 3.3.1) and Lagrangian statistics for the particle phase, considering the effect of Stokes number (§ 3.3.2) and mass loading (§ 3.3.2). Fourier spectra of number density fluctuations are likewise analysed (§ 3.3.3). Finally, conclusions are drawn in § 3.4.

3.2   Methods

3.2.1   Channel flow simulations

The four-way coupled particle-laden channel flow simulations examined in this chapter are a subset of those from chapter 2, as shown in table 3.1. All of the Stokes number variation is included, but only mass loadings up to 100% are considered.

For the statistics reported in the following sections, all simulation cases were run until a steady

	Re_τ	N_p	ρ_p/ρ	ϕ_0	ϕ_{visc}	ϕ_{buf}	ϕ_{log}	St^+	St_η^{log}
Flow 0	229	0	N/A	0	0	0	0	N/A	N/A
Flow 1	225	1.43×10^7	7.42×10^3	0.1	1.37	0.13	0.06	6.90	0.69
Flow 2	227	5.70×10^7	7.42×10^3	0.4	2.36	0.59	0.32	7.01	0.59
Flow 3	234	1.43×10^8	7.42×10^3	1.0	3.54	1.41	0.88	7.45	0.54
Flow 4	263	1.43×10^9	7.42×10^2	1.0	1.52	1.12	0.98	1.02	0.075
Flow 5	273	4.51×10^8	2.35×10^3	1.0	2.63	1.31	0.93	2.98	0.21
Flow 6	213	4.51×10^7	2.35×10^4	1.0	2.79	1.58	0.85	19.6	1.42
Flow 7	210	1.43×10^7	7.42×10^4	1.0	2.10	1.47	0.89	60.0	3.66

TABLE 3.1: Parameters of the channel flow data: the friction Reynolds number Re_τ, the number of particles N_p, the particle to fluid density ratio ρ_p/ρ, the overall mass loading ϕ_0, the average mass loading in the viscous sublayer ϕ_{visc}, the average mass loading in the buffer layer ϕ_{buf}, the average mass loading in the logarithmic layer ϕ_{log}, the friction Stokes number St^+, and the Kolmogorov Stokes number based on the dissipation rate averaged over the logarithmic layer St_η^{log}.

state was reached. Then, simulations were continued for a minimum of 5 flow-through times. Eulerian statistics of time-averaged quantities were obtained by averaging continuously throughout this period, as well as over homogeneous directions. The PDFs of divergence, curl, and helicity were computed from 10 snapshots evenly distributed during the averaging time. Eulerian PDFs were obtained using the same differential operators and grid as the simulation. Lagrangian PDFs were computed with the tessellation technique described in Section 3.2.2.

3.2.2 Tessellation and differential analysis of the particle velocity

Voronoi and Delaunay tessellation

The Delaunay tessellation of the particle positions represents the graph such that no particles are inside the circumscribed sphere of one of any cells. The dual graph of the Delaunay tessellation defines the Voronoi tessellation, which has the property that all points inside are closer to the particle than to the other particles. Figure 3.1(a) shows a particle cluster and the corresponding Voronoi tessellation. We can observe that when the particles are grouped together the Voronoi cell is small, and when the particles are dispersed the Voronoi cell is large. For this reason, small Voronoi cells are characteristic of clusters, and large cells are characteristic of voids. In the following, we use the modified Voronoi tessellation defined in Oujia et al. [153] which uses the centre of gravity of the Delaunay cell to define the vertices of the dual cell C_p, instead of using the circumcentre of the Delaunay cell as done for the Voronoi tessellation. It was shown by Oujia et al. [153] that this construction improves the stability of the numerical method when computing divergence and curl. For computing the 3D tessellation [10], we apply the Quickhull algorithm [18] to the particle positions. In order to take into account the domain walls, the cells corresponding to particles close to the wall and belonging to the convex hull are altered. Each altered cell is composed of the vertices

(a) (b)
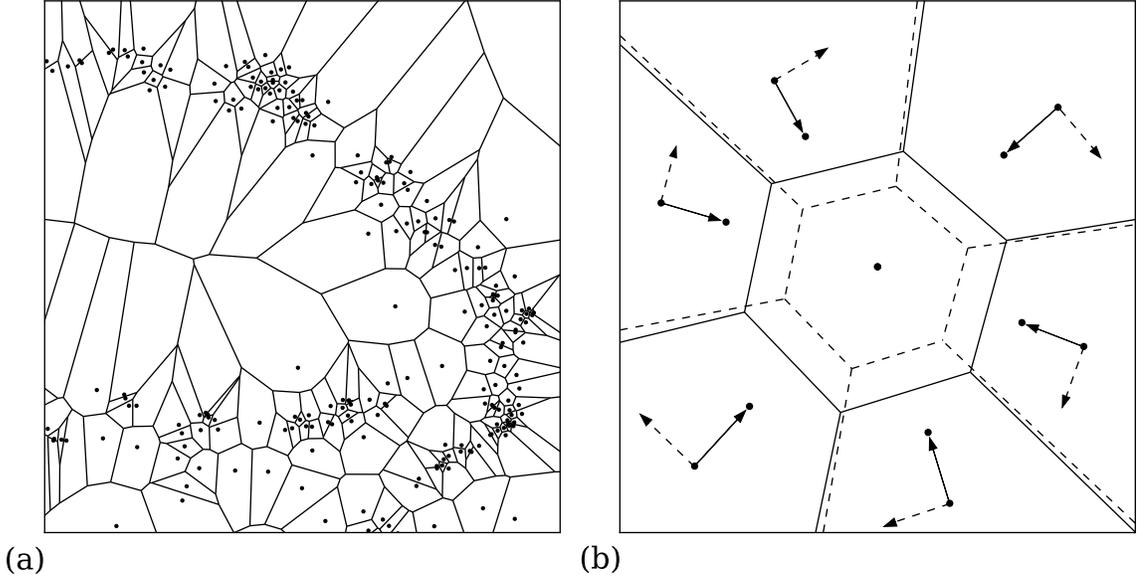
FIGURE 3.1: (a) Voronoi tessellation for a cluster of particles. (b) Example of the motion of a two-dimensional modified Voronoi tessellation at two subsequent time instants for a converging cluster advected by a velocity $-\boldsymbol{v}_{p,z}^{\perp}$. The solid line corresponds to the time step t^k and the dashed line to t^{k+1}. The solid arrows indicate the particle velocity $-\boldsymbol{v}_{p,z}^{\perp}$ and the dashed arrows indicate the particle velocity $\boldsymbol{v}_p$ projected onto the plane orthogonal to the z-direction.

of the cell generated by the modified Voronoi tessellation to which we add the orthogonal projection of these vertices onto the wall. This procedure also prevents degenerated cells, i.e. cells whose particle lies outside its tessellation cell. Note that the number of these modified cells is extremely small: their contribution to the PDFs is of the order of 10^{-6}–10^{-7}.

Divergence, curl, and helicity

To compute the divergence of the particle velocity $\mathcal{D}(\boldsymbol{v}_p)$, following the lines of [151, 153], we first compute the local number density averaged over a cell C_p, which is the inverse of the cell volume. Particles satisfy the conservation equation of the number density n: $D_t n = -n\nabla \cdot \boldsymbol{v}$, where $D_t = \partial_t + \boldsymbol{v} \cdot \nabla$ is the Lagrangian derivative. Considering two subsequent time instants t^k and $t^{k+1} = t^k + \Delta t$ of the modified Voronoi tessellation with time step Δt, we can determine the volume change. Thus we obtain the divergence of the particle velocity as

$$\mathcal{D}(\boldsymbol{v}_p) = -\frac{1}{n}D_t n \approx \frac{1}{2\Delta t}\left(\frac{1}{V_p^{k+1}} + \frac{1}{V_p^k}\right)(V_p^{k+1} - V_p^k), \qquad (3.1)$$

where V_p^k is the modified Voronoi cell volume at t^k. A small timestep is used to ensure that the divergence is calculated accurately. Across all the simulation cases, the timesteps used lie in the

range $4 \times 10^{-4} < \Delta t^+ < 8 \times 10^{-4}$. Figure 3.1(b) shows the dynamics of a modified Voronoi cell at two subsequent time instants. We can observe the decrease of the volume of the cell in the case where the particles are compressed, which corresponds to a convergence of the particle velocity (negative divergence).

Similarly, the curl of the particle velocity can be defined by computing the circulation of the velocity field of particles over a cell C_p. This can also be expressed as the divergence of the velocity which has been projected and rotated in a direction $-\pi/2$ with respect to plane orthogonal to each component of the curl.

Defining $\boldsymbol{v}_x^\perp = \boldsymbol{L}_x\boldsymbol{v}$, $\boldsymbol{v}_y^\perp = \boldsymbol{L}_y\boldsymbol{v}$, and $\boldsymbol{v}_z^\perp = \boldsymbol{L}_z\boldsymbol{v}$ where

$$\boldsymbol{L}_x = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 1 & 0 \end{pmatrix}, \; \boldsymbol{L}_y = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ -1 & 0 & 0 \end{pmatrix}, \text{ and } \boldsymbol{L}_z = \begin{pmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

are rotation matrices around the different axes, the curl of the particle velocity $\mathcal{C}(\boldsymbol{v}_p)$ is obtained by

$$\mathcal{C}(\boldsymbol{v}_p) = \begin{pmatrix} \mathcal{D}\left(-\boldsymbol{v}_{p,x}^\perp\right) \\ \mathcal{D}\left(-\boldsymbol{v}_{p,y}^\perp\right) \\ \mathcal{D}\left(-\boldsymbol{v}_{p,z}^\perp\right) \end{pmatrix}. \tag{3.2}$$

Figure 3.1(b) shows the particle velocity $\boldsymbol{v}_p$ which has been projected onto the plane orthogonal to the z-direction. Following equation (3.2), the z-component of the curl of the particle velocity $\boldsymbol{v}_p$ can be obtained by computing the divergence of $-\boldsymbol{v}_{p,z}^\perp$.

For more details and a thorough validation in the case of a one-way coupled particle-laden isotropic turbulence, we refer to [153]. When computing statistics of curl, the mean shear and symmetry in the wall-normal direction must be accounted for. The z-component of curl, $\mathcal{C}_z$, is dominated by mean shear $(d\bar{u}/dy)$, which changes sign at $y = \delta$. For this reason, in the statistics that follow, $\mathcal{C}_z$ is used in the bottom half of the channel, and its negative is used in the top half. The data are lumped together to compute PDFs.

Having access to the curl of the particle velocity, the helicity of the particle velocity can be computed. Helicity is defined as the scalar product of vorticity and velocity, and yields geometrical information on the alignment of both vector quantities. Geometrical statistics can thus be computed and swirling motion of particle clouds can be quantified. The relative helicity of the particle velocity is defined as

$$\mathcal{H}(\boldsymbol{v}_p) = \frac{\boldsymbol{v}_p \cdot \mathcal{C}(\boldsymbol{v}_p)}{||\boldsymbol{v}_p||_2 \; ||\mathcal{C}(\boldsymbol{v}_p)||_2}, \tag{3.3}$$

and yields the cosine of the angle between the two vectors at each particle position. The range lies between -1 and $+1$, corresponding respectively to anti-alignment and alignment of vorticity and velocity, i.e. strong swirling motion. Two-dimensional motion corresponds to values of 0, reflecting

orthogonality between vorticity and velocity.

3.2.3 Fourier spectra of number density fluctuations

The Fourier spectra of number density fluctuations are calculated for two-dimensional slices in each layer based on the method in [127]. In this method, an analytical Fourier transform is applied to the particle number density in the sliced layer, which is given by

$$n_s(\boldsymbol{x}) = \frac{1}{n_{s0}} \sum_{m=1}^{N_{ps}} \delta_{\text{Dirac}} \left(\boldsymbol{x} - \boldsymbol{x}_p^{(m)} \right), \tag{3.4}$$

where N_{ps} is the number of particles in the sliced layer. The mean number density in the sliced layer is used for the scaling factor n_{s0}. The two-dimensional discrete Fourier transform of $n_s(\boldsymbol{x})$ is then given by

$$\widehat{n_s}(\boldsymbol{k}) = \frac{1}{N_{ps}} \sum_{m=1}^{N_{ps}} \exp\left(-i\boldsymbol{k} \cdot \boldsymbol{x}_p^{(m)} \right), \tag{3.5}$$

where $\boldsymbol{k}$ is the two-dimensional wavenumber vector, i.e., $\boldsymbol{k} = (l_1 \Delta k_x, 0, l_2 \Delta k_z)$ for integers l_1 and l_2, $\Delta k_x \delta = 1/2$, and $\Delta k_z \delta = 3/2$. The number density spectrum $E_n(k)$ is given by

$$E_n(k) = \frac{1}{\Delta k} \sum_{k-\Delta k/2 \leq |\boldsymbol{k}| < k+\Delta k/2} \widehat{\Phi}(\boldsymbol{k}), \tag{3.6}$$

where Δk is the wavenumber interval $(3/2)$. $\widehat{\Phi}(\boldsymbol{k})$ is the spectral density function given by $\widehat{\Phi}(\boldsymbol{k}) = \widehat{n_s}(\boldsymbol{k})\widehat{n_s}^*(\boldsymbol{k})$. By substituting (3.5), $\widehat{\Phi}(\boldsymbol{k})$ is given by

$$\widehat{\Phi}(\boldsymbol{k}) = \left[\frac{1}{N_{ps}} \sum_{m=1}^{N_{ps}} \cos\left(\boldsymbol{k} \cdot \boldsymbol{x}_p^{(m)} \right) \right]^2 + \left[\frac{1}{N_{ps}} \sum_{m=1}^{N_{ps}} \sin\left(\boldsymbol{k} \cdot \boldsymbol{x}_p^{(m)} \right) \right]^2 - \frac{1}{N_{ps}}. \tag{3.7}$$

The third term on the right-hand side of (3.7) is introduced to remove the Poisson noise in the spectrum.

3.3 Results

In this section, we first present Eulerian statistics of the fluid-phase channel flow data and then present the Lagrangian particle statistics. We analyse results of divergence, curl, and helicity of the particle velocity as a function of Stokes number St^+, mass loading ϕ_0, and in different flow regions, i.e. in the viscous sublayer, the buffer layer, and the log layer. In this section, the channel flow is broken up into three regions based on the wall-normal distance. For this purpose, the viscous sublayer is defined as the region $y^+ = [0, 5]$, the buffer layer is defined as $y^+ = [5, 30]$ and the

logarithmic (log) layer is defined as $y^+ = [50, Re_\tau]$. Strictly speaking, this definition includes the wake region as well, but the particle clustering observed in these two regions is very similar.

3.3.1 Eulerian statistics

In four-way coupled turbulent channel flow, particles alter the fluid velocity, both in the mean and fluctuations. In addition, particles tend to accumulate near the channel walls due to turbophoresis [168, 123]. Figure 3.2 shows the degree of these changes in time-averaged quantities, depending on the mass loading and Stokes number. Increasing mass loading decreases the slope of u^+ in the log layer, due to the additional inertia added by particles. Stokes number has less of an effect on the log layer slope, but adds an offset, due to thickening of the viscous sublayer. Figure 3.2(b) shows the local mass loading normalised by the total mass loading, which peaks very close to the wall due to turbophoresis. As mass loading is increased, collisions are more frequent, disrupting the tendency to accumulate near the wall. At lower values of ϕ_0, this results in stronger gradients in mass loading through the viscous sublayer, and a greater difference in mass loading from near-wall to centreline. In terms of turbulence, the fluid phase TKE, shown in 3.2(c), is attenuated slightly as mass loading is increased, and the fluid Reynolds stress, shown in 3.2(d), is attenuated strongly. This is due to an attenuation of the pressure-strain term in the TKE budget, which prevents transfer of energy from the streamwise to other directions, making the Reynolds stress more anisotropic [117]. Within the same mass loading ($\phi_0 = 1$), the effect of increasing Stokes number is to decrease the fluid phase Reynolds stress and increase the fluid-phase TKE up to a point. This is due to the enhancement of streamwise velocity fluctuations and attenuation of wall-normal fluctuations. [114] also observed the attenuation in Reynolds stress as St^+ is increased at $\phi_0 = 0.3$, but not the enhancement in streamwise velocity fluctuations observed here (at $\phi_0 = 1$). Zhao et al. [218] observed enhancement in streamwise velocity fluctuations with increasing St^+, but mass loading was increased at the same time. Therefore, it seems that the degree of enhancement or attenuation in TKE may be dependent on the specific combination of Stokes number and mass loading.

Next, we consider the distribution of the fluid vorticity in different regions of the flow. Figure 3.3 shows PDFs of the three components of the fluid curl in the different flow regions, for the different mass loadings considered, including the no particle case ($\phi_0 = 0$), normalised by the viscous time scale, $\mathcal{C}^+ = \tau_{visc}\mathcal{C}$. This shows the influence of the particle feedback force, which increases in proportion to mass loading. In the viscous sublayer - figure 3.3(a) - the z-component of curl is centred around -1 due to the mean shear which dominates this region. The distribution of $\mathcal{C}_z^+$ is also highly skewed, and events with $\mathcal{C}_z^+ > 0$, i.e. temporary reversal of mean shear, are extremely unlikely. The $\mathcal{C}_x^+$ and $\mathcal{C}_y^+$ distributions are nearly symmetric, as expected from the flow geometry. Moving toward the log layer, the distribution of $\mathcal{C}_z^+$ becomes more symmetric, and the three components more similar, indicating the greater isotropy of turbulence in the log layer. Neither $\mathcal{C}_y^+$ nor $\mathcal{C}_z^+$ is significantly affected by mass loading. However, $\mathcal{C}_x^+$ is strongly attenuated as ϕ_0 increases. While
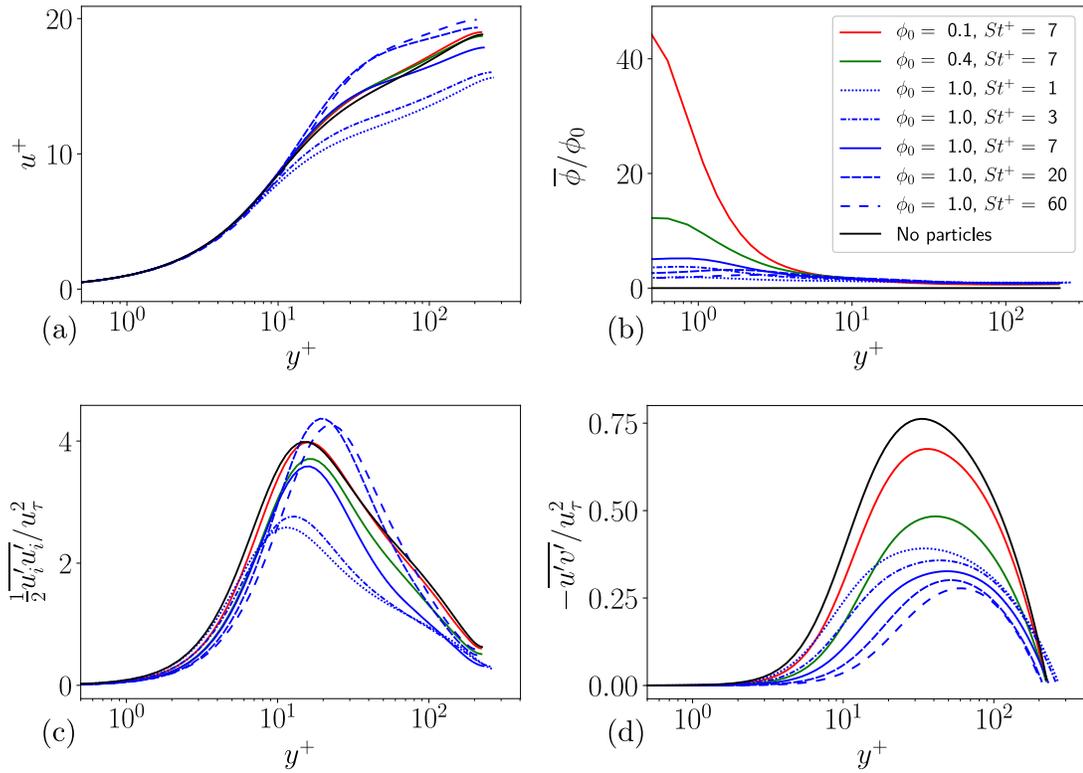
FIGURE 3.2: (a) Mean streamwise fluid velocity, u^+, as a function of wall-normal distance, y^+, for various mass loadings and Stokes numbers, as well as without particles. (b) Local mean mass loading. (c) Fluid phase turbulent kinetic energy. (d) Fluid phase Reynolds shear stress ($\overline{u'v'}$). The legend in (b) applies to all sub-figures.

FIGURE 3.3: PDF of fluid curl, normalised by the viscous time scale, for different mass loadings in the (a) viscous sublayer, (b) buffer layer and (c) log layer. The line colour indicates the component of the curl, and the line style indicates the mass loading.

a small degree of anisotropy is expected in the log layer for low Reynolds number channel flow [7], the extreme attenuation of C_x^+ relative to the other two components is unique to particle-laden wall turbulence. As ϕ_0 is increased, particle-laden wall turbulence becomes more anisotropic, dominated by streamwise velocity fluctuations [117]. These streamwise velocity fluctuations can contribute to vorticity fluctuations in $\mathcal{C}_y^+$ and $\mathcal{C}_z^+$, but not $\mathcal{C}_x^+$, resulting in the selective attenuation.

The effect of St^+ on the PDFs of fluid vorticity is shown in figure 3.4(a-c) for small St^+ and figure 3.4(d-f) for large St^+. For small St^+, the variation in St^+ has little effect on $\mathcal{C}_y^+$, and $\mathcal{C}_x^+$ is attenuated more as St^+ is increased, much like the variation in mass loading.   This can again be explained by the increasing anisotropy of the Reynolds stress with St^+, which was also noted by Lee and Lee [114]. $\mathcal{C}_z^+$ variance is somewhat enhanced in the log layer, perhaps due to the stronger streamwise velocity fluctuations with increasing St^+. As St^+ increases beyond $St^+ = 7$, all three curl components are attenuated as these trends continue, and $\mathcal{C}_z^+$ is is preferentially attenuated at negative values, resulting in a narrower, less skewed distribution as particle motions (and the drag forces particles exert) become more random, and less determined by eddies. The effect of this change on the fluid phase is to weaken rotational motion (i.e. eddies) at large Stokes numbers.

Finally, the fluid-phase relative helicity distributions are plotted in figure 3.5 for the effect of ϕ_0 and figure 3.6 for the effect of St^+. The helicity distributions are all symmetric with zero mean due to the problem geometry, but they are strongly affected by flow region. The probability of having zero helicity is highest in the viscous sublayer for all cases. This is expected because the near-wall vorticity is dominated by the mean shear ($d\bar{u}/dy$), which is orthogonal to the streamwise velocity. In other words, in the viscous sublayer and buffer layer there is an absence of swirling flow features, and the flow is nearly two-dimensional. In the buffer layer, the helicity distribution is less strongly peaked, and in the log layer, the distribution is nearly uniform. Across all three layers, the effect of increasing the mass loading or Stokes number is to make the flow more two-dimensional. In the
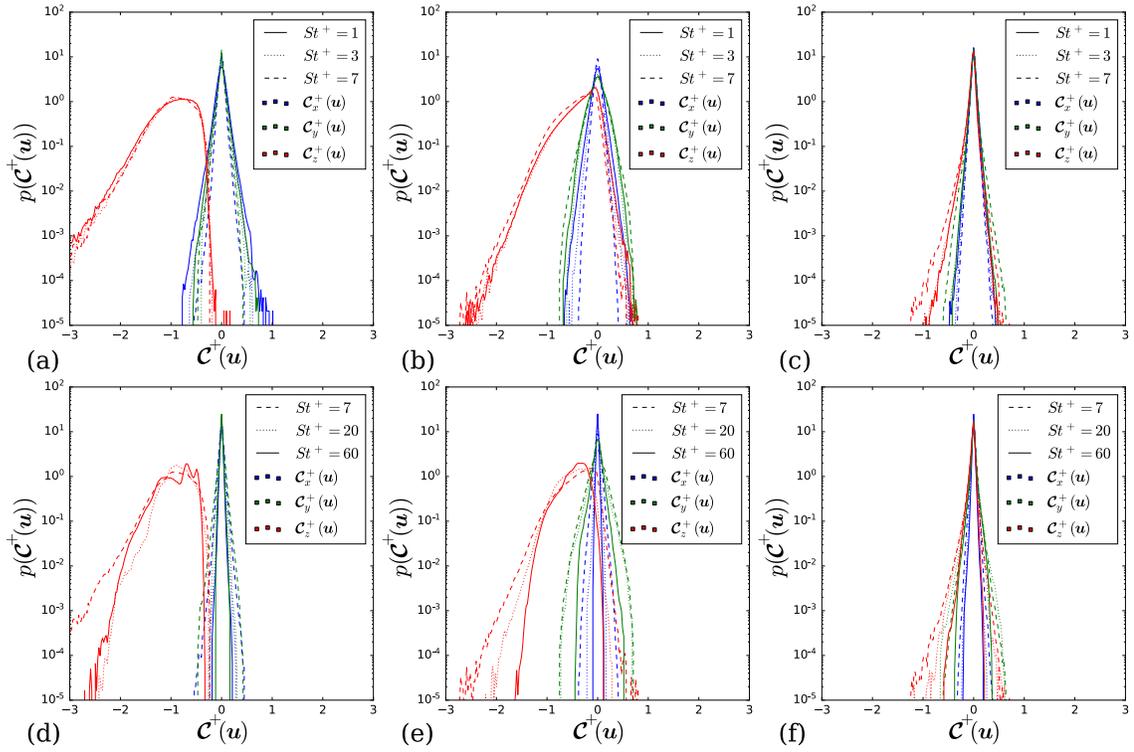
FIGURE 3.4: PDF of fluid curl, normalised by the viscous time scale, for a range of Stokes numbers. Data for small Stokes numbers: (a) viscous sublayer, (b) buffer layer and (c) log layer. Data for large Stokes numbers: (d) viscous sublayer, (e) buffer layer and (f) log layer. The line colour indicates the component of the curl, and the line style indicates the Stokes number.
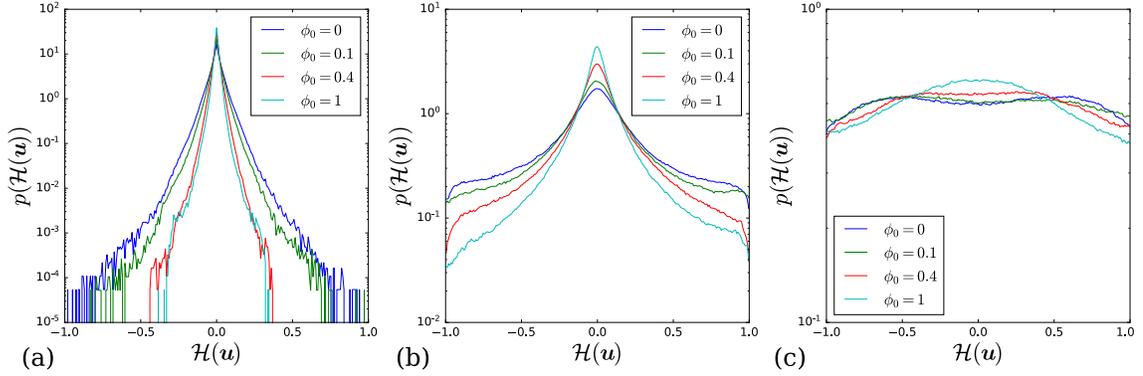
FIGURE 3.5: PDF of fluid helicity for various mass loadings in the (a) viscous sublayer, (b) buffer layer and (c) log layer. Note that the y-axis ranges for each PDF are different.
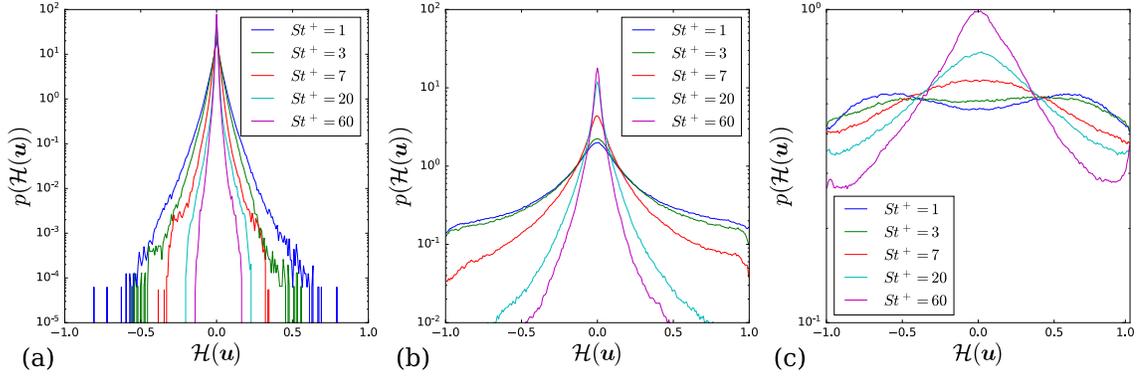
FIGURE 3.6: PDF of fluid helicity for various Stokes numbers in the (a) viscous sublayer, (b) buffer layer and (c) log layer. Note that the y-axis ranges for each PDF are different.

viscous sublayer and buffer layer, this can be interpreted as a consequence of the low-speed streaks becoming more stable and persistent, due to two-way coupling. In the log layer, for sufficiently small ϕ_0 and St^+, there are two peaks at moderate values of helicity. Previous observations from HIT [152] have shown peaks at $\mathcal{H}_p = \pm 1$ for small Stokes number, indicating a greater probability of swirl than non-swirl. The peaks at moderate values observed here may be due to the increased anisotropy of particle-laden turbulence, which limits how three-dimensional its coherent structures can be. We can conjecture that the curves obtained at $St^+ = 1$ and $\phi_0 = 0.1$ represent a transition from a concave to convex PDF, and that if St^+ and ϕ_0 were decreased further, the curves would become completely convex as observed for the HIT case in Oujia et al. [152].

3.3.2 Lagrangian statistics

Influence of Stokes number

In this section we consider the effect of Stokes number on cluster dynamics by examining the particle velocity divergence, curl, and helicity for $St^+ = [1-60]$. It is important to underscore that all cases considered in this section have the same mass loading, $\phi_0 = 1$, at which strong two-way coupling between the phases, as well as significant changes in the turbulence anisotropy, are expected, as described in Section 3.3.1. We plot particle positions coloured by particle velocity divergence in figure 3.7 for $St^+ = 3$ and $St^+ = 60$. For $St^+ = 3$, shown in figure 3.7(a-c) the particles in the viscous sublayer organise into elongated chain-like structures, which correspond to the low-speed streaks. The effect of particles on the low speed streaks, as well as other flow structures, is shown with visualisations in Appendix 3.5. Clusters of particles with large positive and negative divergence are sparsely distributed, but the majority of particles have near-zero divergence, indicating that most particles are simply transported by the flow (indicated by green colour). In the clusters with non-zero divergence, positive and negative divergence are intermixed, suggesting crossing trajectories and collisions. Thin, elongated void regions, indicated by white space, appear between the particle streaks. In the buffer layer, clusters of particles with negative and positive divergence are more densely distributed, but they are less elongated. The more intense convergence/divergence in the buffer layer is likely due to the increased strength of velocity fluctuations there, as shown in figure 3.2(c). In the log layer, the density of clusters with non-zero divergence is decreased, and their structures are less directional, bearing a strong resemblance to previous results in HIT for small Stokes numbers [151]. For $St^+ = 60$, shown in figure 3.7(d-f), the particle distributions and their velocity divergence are starkly different. The width of the streaky structures has greatly increased, matching the structure of the streamwise velocity field (Appendix 3.5, figure 3.23(d,e)). Throughout the layers, a much greater proportion of particles have non-zero divergence, indicating that particles are moving more independently of the surrounding fluid. Finally, we observe preferential concentration in the log layer for $St^+ = 60$, but only at large scales, as expected because of large particle inertia.

Figures 3.8 and 3.9 show snapshots of particles coloured by the magnitude of curl and helicity, respectively, using the same snapshots as in figure 3.7. In the viscous sublayer, this shows that the low-speed streaks where the divergence is near-zero correspond to low vorticity magnitude, due to the weaker shear in low speed streaks. Furthermore, the large values of divergence correspond to the regions of large vorticity magnitude. This trend applies to both small and large Stokes numbers, and continues into the buffer and log layers, suggesting that the rotational motion of particle clouds is responsible for divergence and convergence as particles are centrifuged out of eddies. The vortical motion of particles in the log layer is similar to what has been observed for HIT in Oujia et al. [152].

The helicity of the particle velocity in figure 3.9 illustrates that particle clusters in the log layer exhibit a greater alignment and anti-alignment of particle velocity and rotation, corresponding
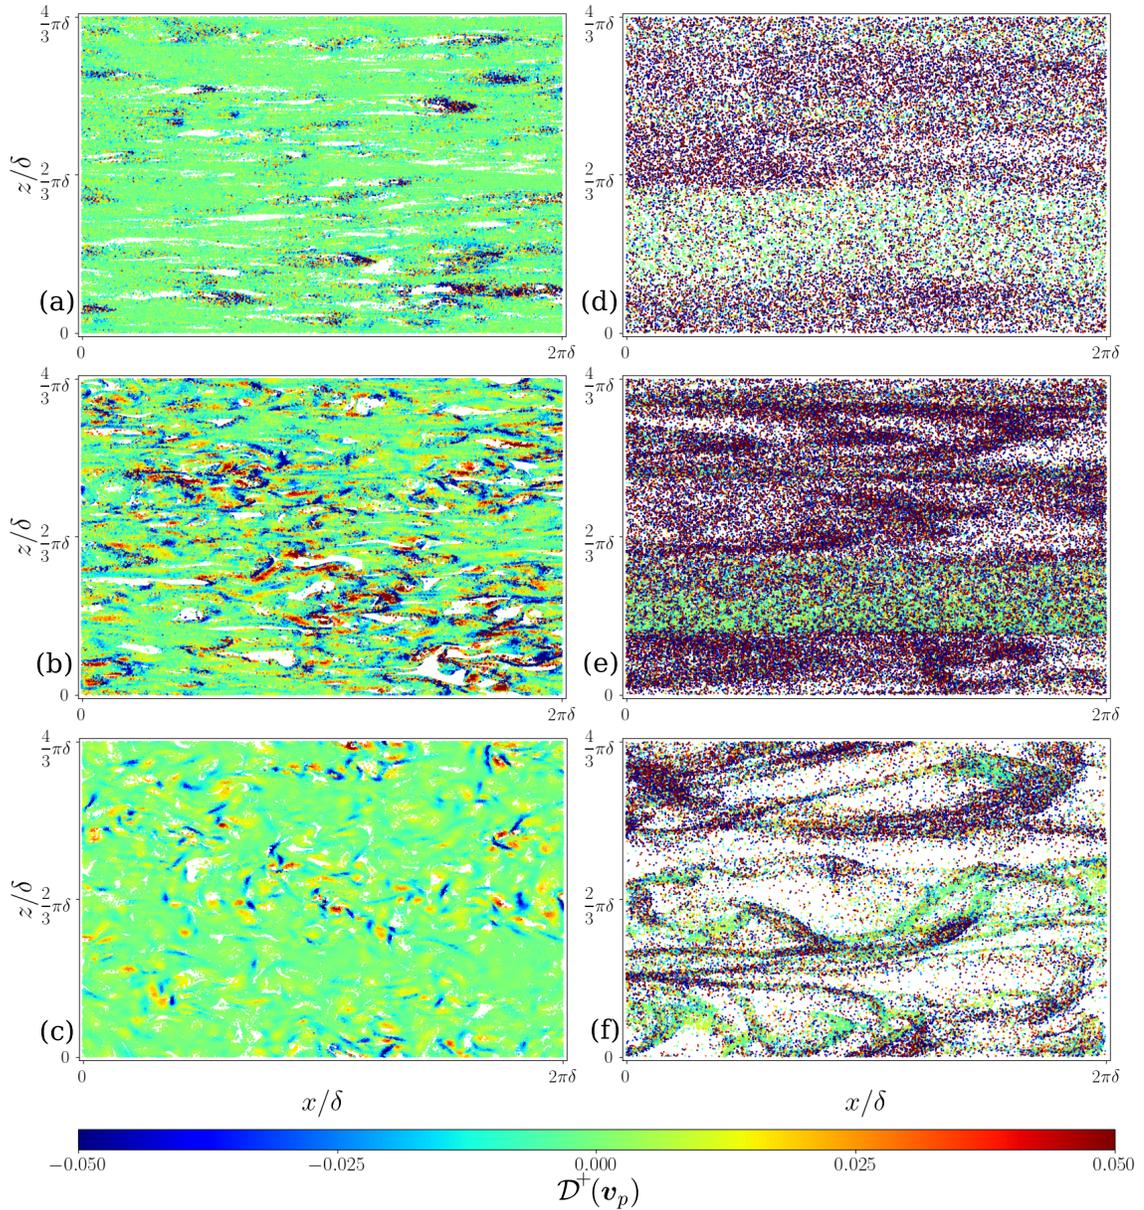
FIGURE 3.7: Particle positions coloured by divergence in wall-parallel slices, for (a) viscous sublayer, $St^+ = 3$, $\phi_0 = 1$; (b) buffer layer, $St^+ = 3$, $\phi_0 = 1$; (c) log layer, $St^+ = 3$, $\phi_0 = 1$; (d) viscous sublayer, $St^+ = 60$, $\phi_0 = 1$; (e) buffer layer, $St^+ = 60$, $\phi_0 = 1$; (f) log layer, $St^+ = 60$, $\phi_0 = 1$. Divergence is normalised by the viscous time scale, and the colour bar is truncated to emphasise regions of zero divergence. Because of the difference in number density, thinner slices are used to visualise the $St^+ = 3$ data. The slices used for $St^+ = 3$ are $2.375 \leq y^+ < 2.625$, $11.375 \leq y^+ < 12.625$ and $99.375 \leq y^+ < 100.625$ for viscous sublayer, buffer layer and log layer, respectively. The slices used for $St^+ = 60$ are $2 \leq y^+ < 3$, $9.5 \leq y^+ < 14.5$ and $97.5 \leq y^+ < 102.5$ for viscous sublayer, buffer layer and log layer, respectively.

to dark blue and dark red particle regions, respectively. This confirms that particle clusters are swirling, i.e. their motion is helical. In contrast, in the viscous sublayer and the buffer layer helical motion is less present, since vorticity and velocity are more perpendicular there, resulting in two-dimensional motion of the particles. Stronger swirl is observed in figure 3.9(c) compared to figure 3.9(f), indicating that the particle motion becomes more two-dimensional as Stokes number is increased, even in the log layer.

To characterise the presence of clusters and voids, figure 3.10 shows the PDFs of modified Voronoi volumes for each flow region and Stokes number normalised by the mean modified Voronoi volume within each layer $\overline{V}_{p,l}$. This is done because the particle number density is different within each layer due to turbophoresis, so this normalization facilitates comparison with a random distribution within each layer. Across all flow regions, there is clear deviation from a random particle distribution with greater probabilities of both clusters and voids (small and large values of $V_{p,l}/\overline{V}_{p,l}$, respectively). At both very small and very large St^+, the distribution is closer to random than at intermediate St^+, where preferential concentration plays the greatest role. In general, the distribution of clusters in the viscous sublayer is closer to random than in other regions of the flow, except for $St^+ = 1$, in which the clusters are very close to random throughout the channel. This is due to the organisation into the low speed streaks, within which particles remain for a long time and become evenly distributed (figure 3.7(a)). The large modified Voronoi volumes, i.e. voids, depart significantly from random in all layers, but especially in the viscous sublayer, where large, elongated voids between the streaks are present. The PDFs become noisy at their right tails due to a lack of samples. The distribution in the viscous sublayer is the noisiest because of the normalisation used. Particles are more densely packed in the viscous sublayer due to turbophoresis, so the average volume in the layer is smaller, shifting the curve to the right relative to other layers.

Figure 3.11 shows the PDFs of particle velocity divergence for the different Stokes numbers, and table 3.2 gives the variance and flatness values. The divergence data is normalised by the viscous timescale: $\mathcal{D}^+ = \tau_{visc}\mathcal{D}$. As St^+ is increased from a small value, the variance of the particle velocity divergence increases across flow regions, and the tails of the PDFs become heavier, as observed in HIT [151]. At the tails, the PDFs become noisy, due to the rarity of particles with extreme values of divergence, which could be improved by averaging over additional flow snapshots. In the buffer layer, the variance is greatest, due to the strength of turbulent velocity fluctuations. This corresponds to the maximum in turbulent kinetic energy (TKE), which occurs in the range $10 \lesssim y^+ \lesssim 30$ for the flows considered here (see figure 3.2(c)). In the log layer and to a lesser extent the buffer layer, these changes eventually saturate: the increase in variance slows down and the flatness begins to decrease, because the TKE also saturates. This saturation occurs around $St^+ = 7$, which corresponds to $St^{\log}_\eta = 0.54$. Results from HIT show a similar transition at $St_\eta = 1$ [151], which is remarkably close considering the differences in two-way coupling, turbulence anisotropy, and Reynolds number. By contrast, the variance of divergence continues to increase in the viscous sublayer, likely due to the
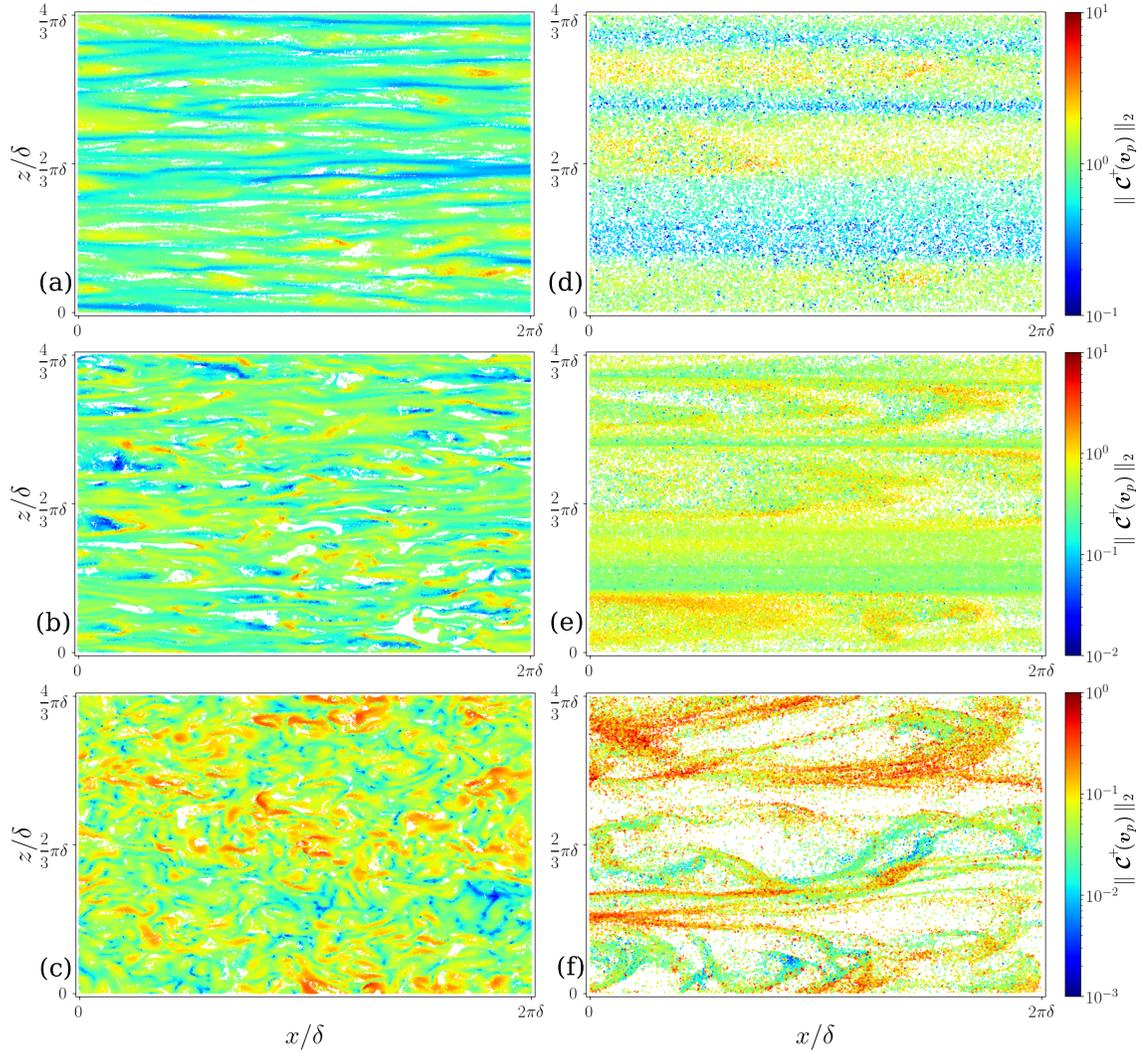
FIGURE 3.8: Particle positions coloured by vorticity magnitude in wall-parallel slices, for (a) viscous sublayer, $St^+ = 3$, $\phi_0 = 1$; (b) buffer layer, $St^+ = 3$, $\phi_0 = 1$; (c) log layer, $St^+ = 3$, $\phi_0 = 1$; (d) viscous sublayer, $St^+ = 60$, $\phi_0 = 1$; (e) buffer layer, $St^+ = 60$, $\phi_0 = 1$; (f) log layer, $St^+ = 60$, $\phi_0 = 1$. Thickness of the layers used are defined in figure 3.7.
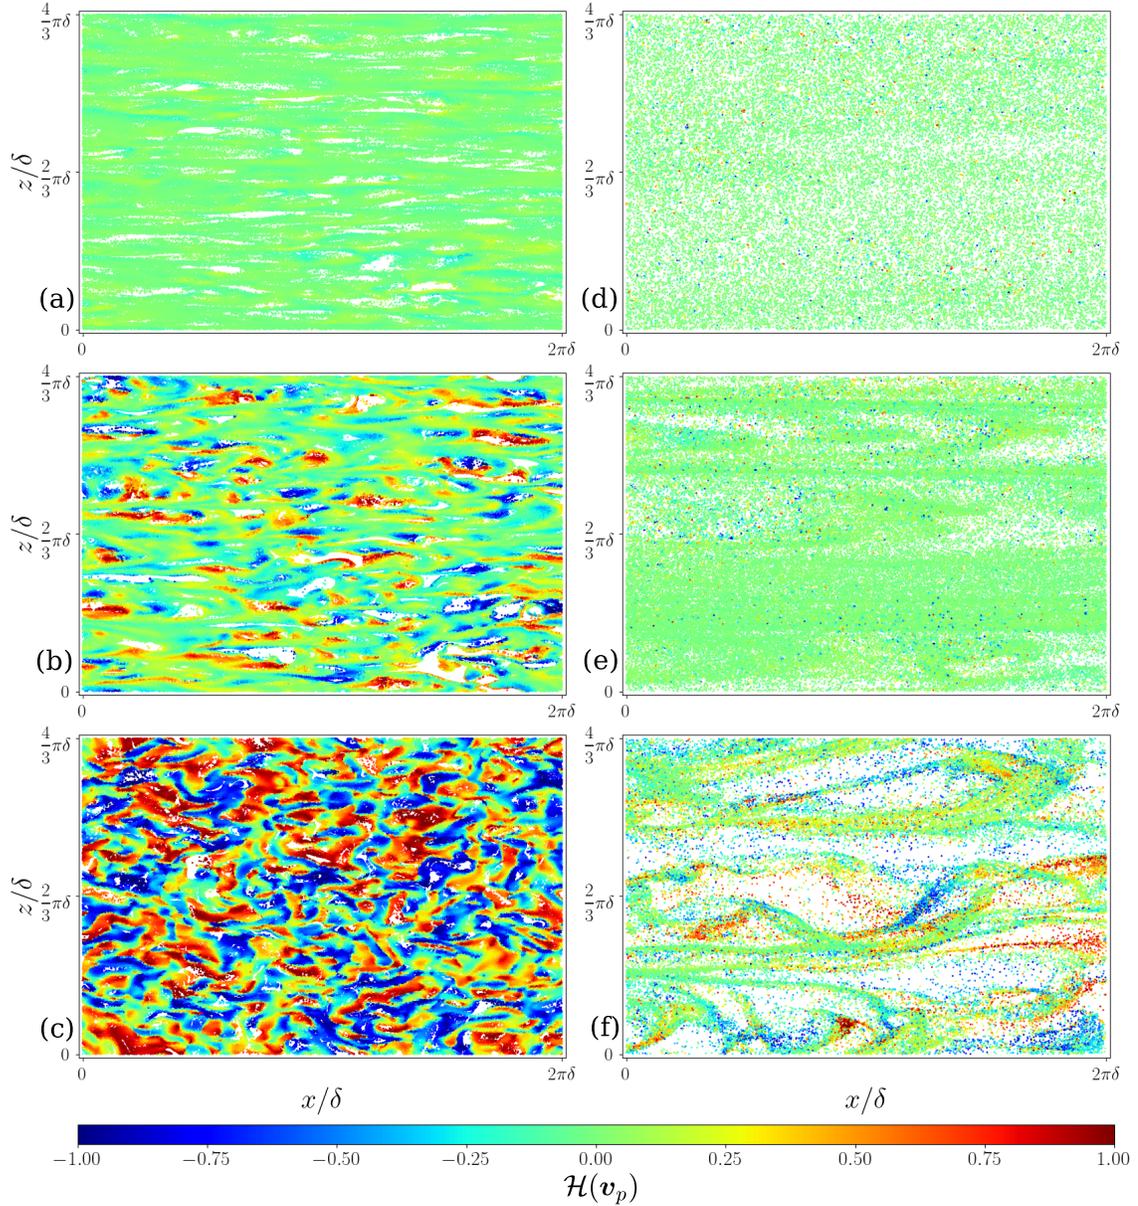
FIGURE 3.9: Particle positions coloured by relative helicity in wall-parallel slices, for (a) viscous sublayer, $St^+ = 3$, $\phi_0 = 1$; (b) buffer layer, $St^+ = 3$, $\phi_0 = 1$; (c) log layer, $St^+ = 3$, $\phi_0 = 1$; (d) viscous sublayer, $St^+ = 60$, $\phi_0 = 1$; (e) buffer layer, $St^+ = 60$, $\phi_0 = 1$; (f) log layer, $St^+ = 60$, $\phi_0 = 1$. Thickness of the layers used are defined in figure 3.7.
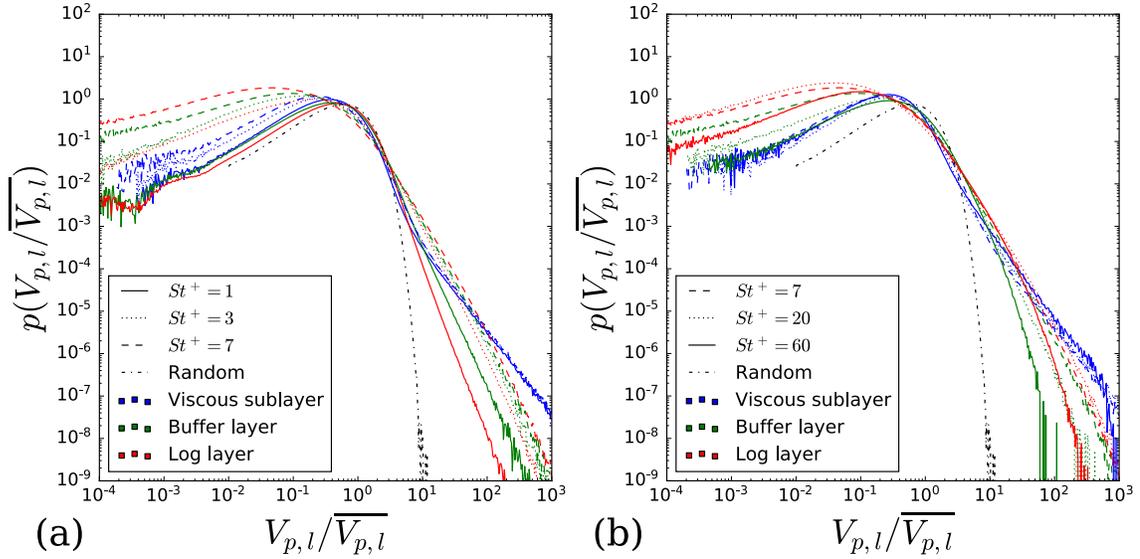
FIGURE 3.10: PDF of modified Voronoi volume as a function of flow region and Stokes number, normalised by the mean modified Voronoi volume within each layer for (a) $St^+ = 1, 3, 7$ and (b) $St^+ = 7, 20, 60$. The dashed-dotted line represents the volume distribution for randomly distributed particles.

increasingly ballistic motion of very heavy particles.

In their analysis of particle clustering, Esmaily and Mani [56] present PDFs of the the exponential particle contraction rate, which is closely related to the divergence of the particle velocity, as a function of Stokes number in one-way coupled homogeneous isotropic turbulence. They show that the variance of this PDF initially increases with Stokes number, like the above results, but that its variance decreases for $St_\eta > 1$. The main difference in their analysis is that they average the contraction rate along Lagrangian trajectories, whereas the present study uses a more Eulerian perspective, considering the divergence of all the particles in their present location, regardless of history. These two perspectives may be reconcilable, and a promising indication to this end is that Esmaily and Mani [56] show that for $St_\eta > 1$, the variance is less sensitive to St_η, which is certainly true in the present results (figure 3.11), where saturation in the particle velocity divergence is observed.

Next, we address the particle velocity curl to characterise the vortical motion of clusters. Figures 3.12(a-c) and 3.12(d-f) show PDFs of the three components of the curl of particle velocity in the different flow regions for small and large Stokes numbers, respectively. Like the divergence, curl has been normalised by the viscous timescale: $\mathcal{C}^+ = \tau_{visc}\mathcal{C}$. The particle-phase PDFs have heavier tails than their fluid counterparts in figure 3.4, indicating greater probability of extreme values, but they show similar trends in terms of asymmetry and anisotropy, due to two-way coupling. Similar
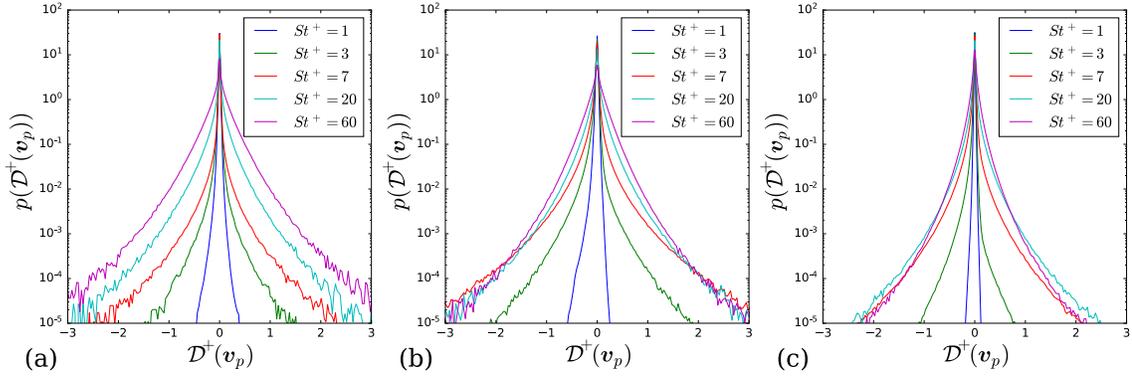
FIGURE 3.11: PDF of particle velocity divergence, normalised by the viscous time scale, for various Stokes numbers in the (a) viscous sublayer, (b) buffer layer, and (c) log layer.

St^+	Viscous	Buffer	Log	St^+	Viscous	Buffer	Log
1	7.93×10^{-5}	2.25×10^{-4}	2.23×10^{-5}	1	1.26×10^2	0.44×10^2	0.30×10^2
3	4.26×10^{-4}	1.88×10^{-3}	2.10×10^{-4}	3	9.91×10^2	2.73×10^2	8.39×10^2
7	1.79×10^{-3}	9.78×10^{-3}	2.47×10^{-3}	7	11.8×10^2	2.77×10^2	3.49×10^2
20	1.23×10^{-2}	1.91×10^{-2}	6.04×10^{-3}	20	0.77×10^2	0.51×10^2	1.07×10^2
60	4.35×10^{-2}	3.25×10^{-2}	7.06×10^{-3}	60	0.18×10^2	0.17×10^2	0.33×10^2

TABLE 3.2: Variance (left) and flatness (right) of the particle velocity divergence normalised by the viscous time scale for different layers and Stokes numbers St^+.
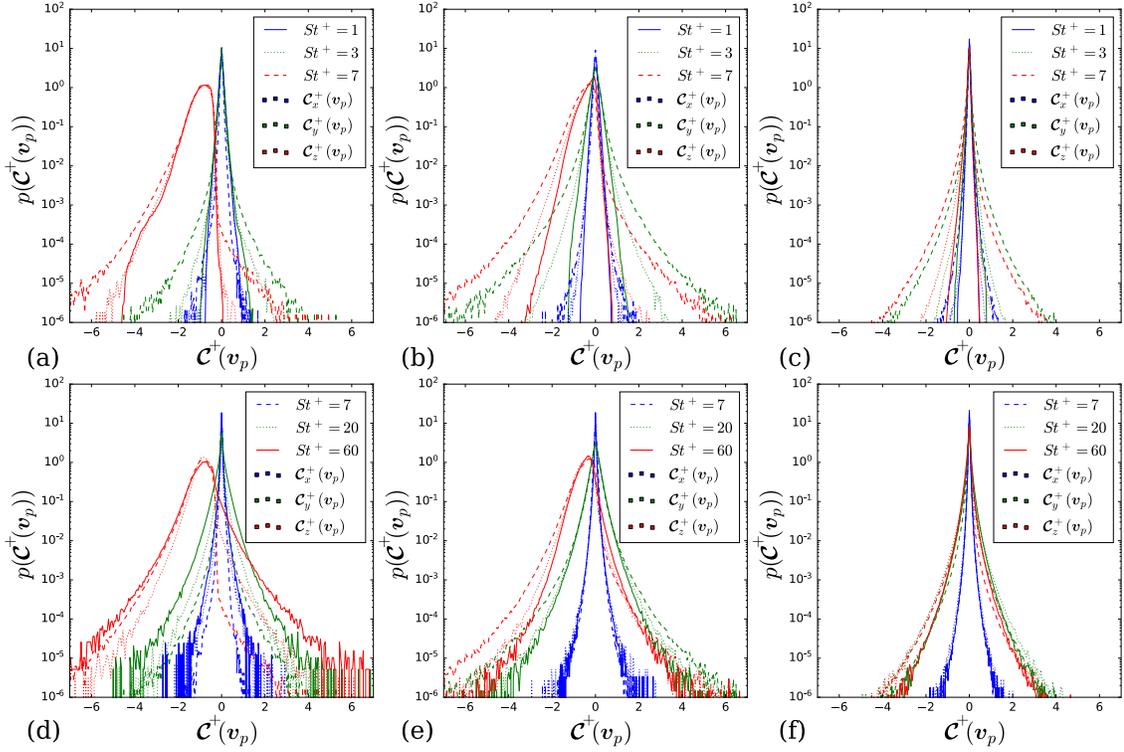
FIGURE 3.12: PDF of particle velocity curl, normalised by the viscous time scale, for a range of Stokes numbers. Data for small Stokes numbers: (a) viscous sublayer, (b) buffer layer and (c) log layer. Data for large Stokes numbers: (d) viscous sublayer, (e) buffer layer and (f) log layer. The line colour indicates the component of the curl, and the line style indicates the Stokes number.

to divergence, the variance of particle velocity curl initially increases as St^+ increases, but then saturates in the log layer for larger St^+. This can be justified because intense fluctuations of fluid vorticity are short-lived, and the large particle inertia prevents the particles from following the streamlines of intense eddies. In the viscous sublayer, the effect on $\mathcal{C}_z^+$ is particularly striking. As St^+ is increased, the distribution of $\mathcal{C}_z^+$ remains offset with a mean value of -1, but it becomes more symmetrical, with significant distributions at $\mathcal{C}_z^+ > 0$. This broadening of the PDF can be explained by the enhancement of the ballistic motions of inertial particles; with increasing inertia, the trajectories tend to become more independent of the streamlines of the carrier fluid as St^+ increases.

Figure 3.13 shows the PDFs of the particle helicity for various Stokes numbers. Similar to the fluid helicity, the particle helicity indicates that the motions of the particles become more two-dimensional approaching the wall. In the log layer, the distribution at large St^+ is concave, and it flattens as St^+ is decreased. The peaks at moderate values of helicity observed for $St^+ = 1$ suggest that the PDF may be transitioning from concave to convex, as observed in Oujia et al. [152].
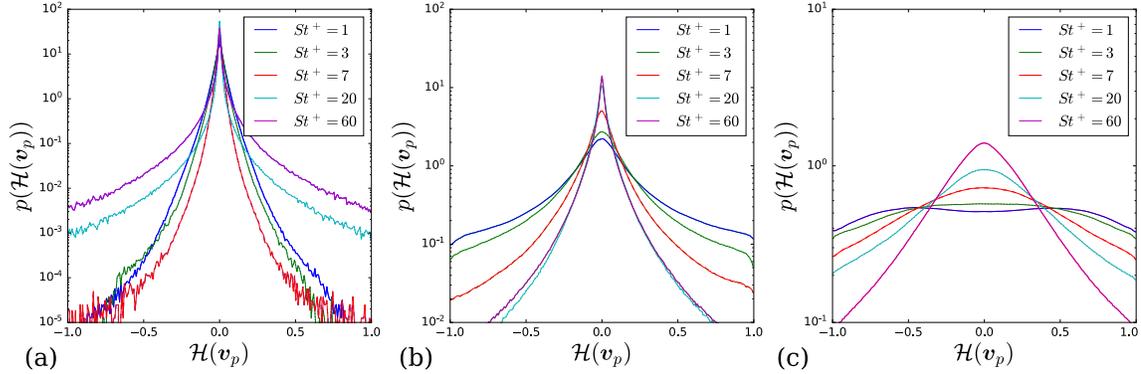
FIGURE 3.13: PDFs of the relative helicity (cosine of the angle between particle velocity and vorticity) $\mathcal{H}(\boldsymbol{v}_p)$ of the particles for various Stokes numbers in the (a) viscous sublayer, (b) buffer layer, and (c) log layer. Note that the y-ranges differ by orders of magnitude.

Interestingly, in the viscous sublayer, a non-monotonic trend is observed, in contrast to the fluid helicity (figure 3.6). Initially the helicity distribution narrows as St^+ increases, but then it suddenly broadens, indicating that the particle motion is becoming much more three-dimensional, and independent of the fluid phase. This would be expected if the particle phase is transitioning away from organised, but attenuated turbulent motion, to ballistic motion driven by particle inertia and collisions.

Influence of mass loading

Next, we examine the effect of mass loading on particle phase Lagrangian statistics. To gain a qualitative understanding of clustering and particle velocity divergence in the different flow regions and for different mass loadings, Figure 3.14 shows particle positions coloured by divergence in wall-parallel slices. In the viscous sublayer, particles are organised into low-speed streaks, and most particles show zero divergence, meaning they are being transported without clustering, similar to the particles in figure 3.7(a), but the streaks are more distinct because the number density is lower. In the buffer layer, the streaky structures are still visible, but the prevalence of convergence and divergence is greatly increased, and it is concentrated in hairpin-like structures. Finally, in the log layer, the clustering becomes similar to that of moderate Stokes number particles in HIT [151]. At larger mass loading, there is a greater proportion of zero-divergence particles, particularly in the viscous sublayer and buffer layers. As ϕ_0 is increased, the particles are more densely packed, but the structures remain qualitatively similar. Only $\phi_0 = 0.1, 0.4$ are shown, but the same qualitative trend continues to $\phi_0 = 1$.

Figure 3.15 shows PDFs of modified Voronoi volumes for each flow region and mass loading. As with the variation in Stokes number, the distribution of clusters and voids in the viscous sublayer comes closer to a random distribution than other flow regions. As the mass loading increases, the
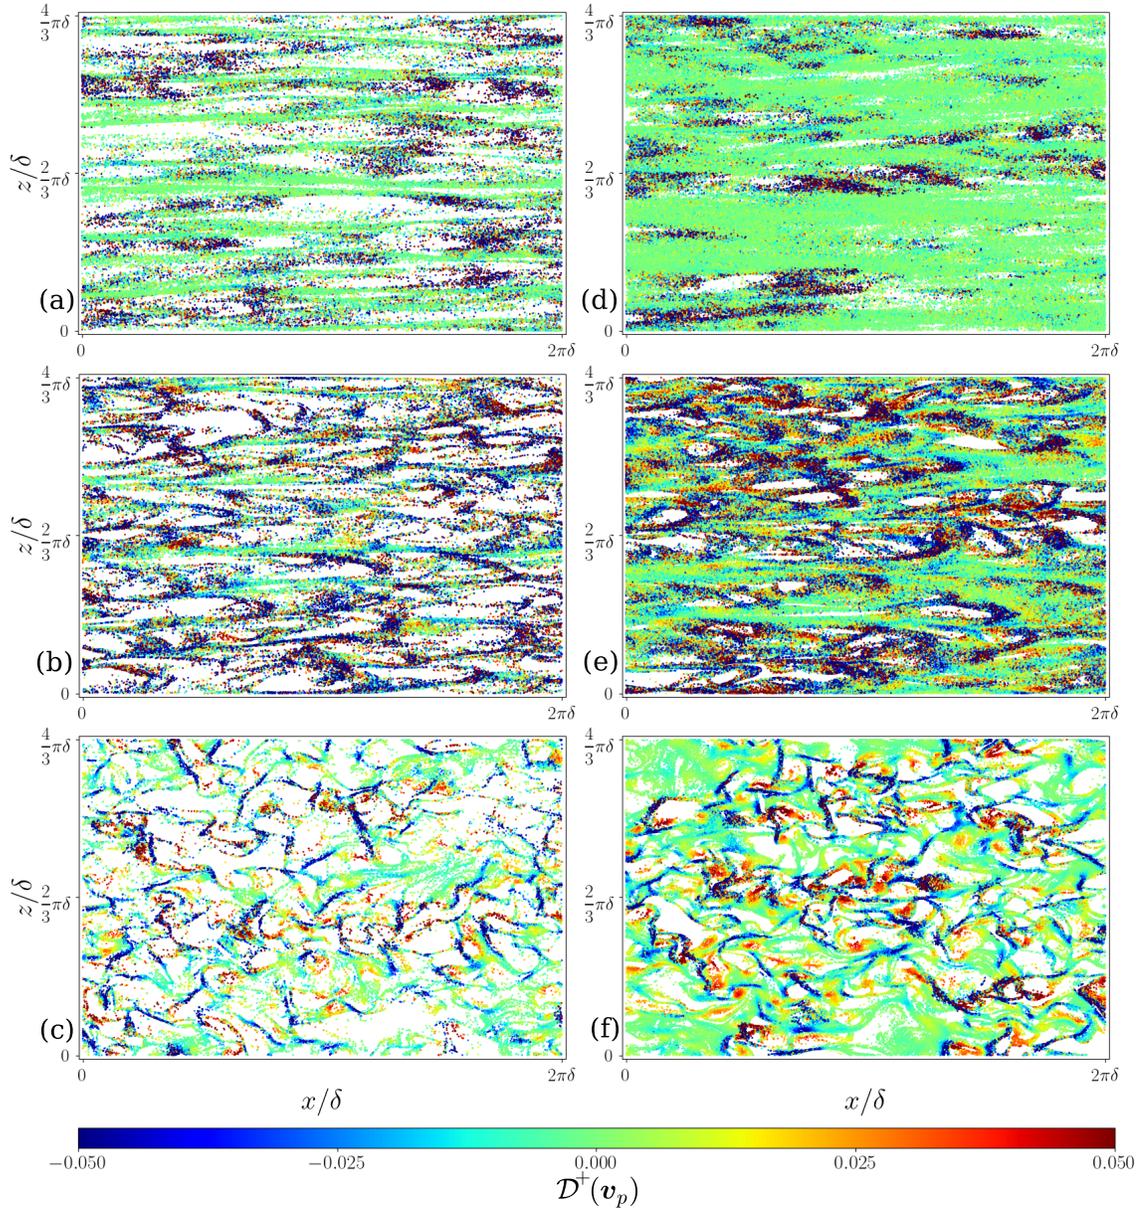
FIGURE 3.14: Particle positions coloured by divergence in wall-parallel slices, for (a) viscous sublayer, $\phi_0 = 0.1$; (b) buffer layer, $\phi_0 = 0.1$; (c) log layer, $\phi_0 = 0.1$; (d) viscous sublayer, $\phi_0 = 0.4$; (e) buffer layer, $\phi_0 = 0.4$; (f) log layer, $\phi_0 = 0.4$. Divergence is normalised by the viscous time scale, and the colour bar is truncated to emphasise regions of zero divergence. The slices used for visualisation are $2 \leq y^+ < 3$, $9.5 \leq y^+ < 14.5$ and $97.5 \leq y^+ < 102.5$ for viscous sublayer, buffer layer, and log layer, respectively.
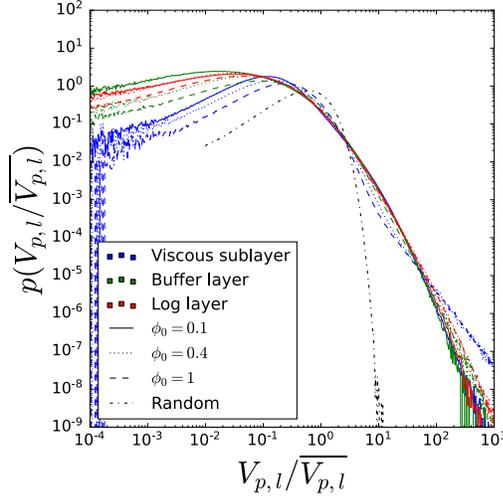
FIGURE 3.15: PDF of modified Voronoi volume as a function of flow region and mass loading, normalised by the mean modified Voronoi volume within each layer. The dashed-dotted line represents the volume distribution for randomly distributed particles.

ϕ_0	Viscous	Buffer	Log	ϕ_0	Viscous	Buffer	Log
0.1	6.06×10^{-3}	10.96×10^{-3}	2.15×10^{-3}	0.1	0.70×10^2	0.56×10^2	1.21×10^2
0.4	2.45×10^{-3}	12.12×10^{-3}	2.54×10^{-3}	0.4	7.21×10^2	1.15×10^2	3.91×10^2
1.0	1.79×10^{-3}	9.78×10^{-3}	2.47×10^{-3}	1.0	11.8×10^2	2.77×10^2	3.49×10^2

TABLE 3.3: Variance (left) and flatness (right) of the particle velocity divergence normalised by the viscous time scale for different layers and mass loadings ϕ_0.

clusters (small volumes) across flow regions become closer to the random distribution. This may be due to increased collisions due to greater number density of particles at higher mass loading. It may also be due to the attenuation of the carrier turbulence causing slightly less preferential concentration. The distribution of voids is strikingly similar as ϕ_0 is increased in the buffer and log layers. This may be because the relatively low mass loading of particles in voids means that two-way coupling does not play a role in altering those structures.

To highlight the effect of mass loading in the different flow regions on particle velocity divergence, we show PDFs of divergence for each region in figure 3.16, and report variance and flatness in table 3.3. The variance is largest in the buffer layer across all flow cases, as expected because the near-wall peak of the fluid turbulent kinetic energy occurs in this region. In the buffer and log layers, the tails of the PDFs become slightly heavier as mass loading is increased. This may be due in part to the larger number of particles as ϕ_0 is increased, which offer more sampling to capture strong convergence/divergence, which occur less frequently, as described by [151]. This issue affects the viscous sublayer least, because turbophoresis makes the number density similar there, despite the
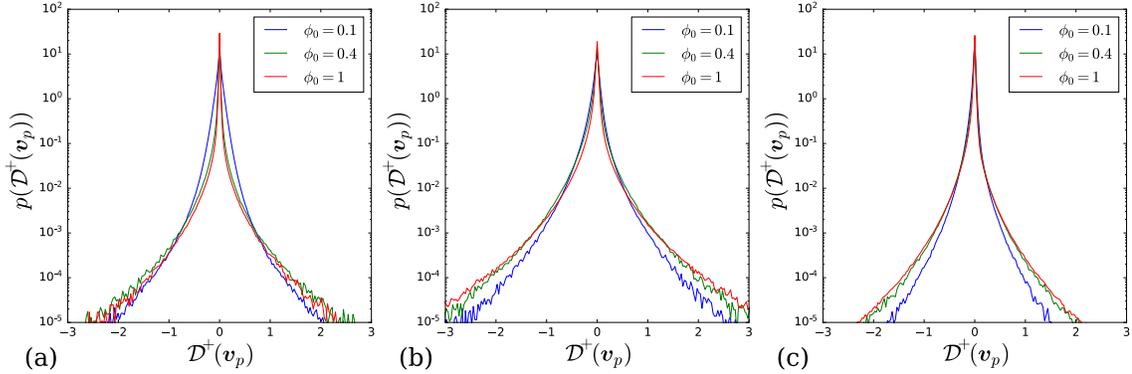
FIGURE 3.16: PDF of particle velocity divergence, normalised by the viscous time scale, for three different mass loadings in the (a) viscous sublayer, (b) buffer layer, and (c) log layer.

overall change in mass loading. The variance of the divergence in the viscous sublayer decreases as ϕ_0 increases, which is likely due to the greater fluid turbulent kinetic energy at low mass loading. The greater variance for $\phi_0 = 0.1$ can be seen in figure 3.14(a) and 3.14(d), where a greater proportion of particles have zero divergence (green) at the higher mass loading. The changes to particle velocity divergence in the buffer and log layer due to mass loading are less substantial, though there may be a weak maximum in variance at $\phi_0 = 0.4$, as shown in table 3.3.

Figure 3.17 shows the three components of particle velocity curl across the flow regions. Like the previously shown PDFs of curl, $\mathcal{C}_z^+$ is asymmetric with a negative mean, while $\mathcal{C}_x^+$ and $\mathcal{C}_y^+$ are nearly symmetric. Heavier tails as ϕ_0 is increased are again observed, possibly due to increased sampling at larger ϕ_0. A significant change due to mass loading is observed in the viscous sublayer, where the larger mass loading shows a narrower PDF with strongly decreased likelihood of $\mathcal{C}_z^+ > 0$, i.e. locally reversed flow. Narrowing of the distribution is expected based on the tendency of particles to stabilise the near-wall streaks and attenuate velocity fluctuations. Similar to divergence, the particle velocity curl in x and y directions has the greatest variance in the buffer layer, which suggests that active formation and annihilation of clusters is associated with stronger vortical motions. As in the fluid vorticity PDFs shown in figure 3.3, $\mathcal{C}_x^+$ has a smaller variance than $\mathcal{C}_y^+$ and $\mathcal{C}_z^+$. Both fluid and particle curl show evidence of anisotropy due to two-way coupling. The fact that the curl is most isotropic in the log layer is also reflected in the particle visualisations in figure 3.14(c,f), which are qualitatively similar to previous results from HIT [151].

Figure 3.18 shows PDFs of the relative helicity of particle velocity for different mass loadings in the different layers. Much like the other results, the particle helicity is more likely to be zero (two-dimensional flow) the closer particles get to the wall. The effect of mass loading in the buffer layer and log layer is small, but in the viscous sublayer, the particle helicity distribution is significantly flatter with $\phi_0 = 0.1$. This indicates that the particle motion is much more three dimensional in this case. While it might be expected that this is due to a smaller degree of two-way coupling, this
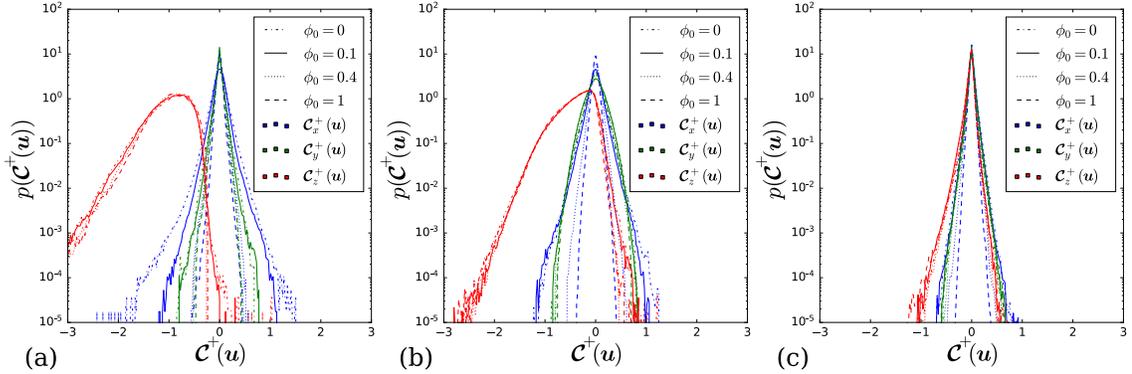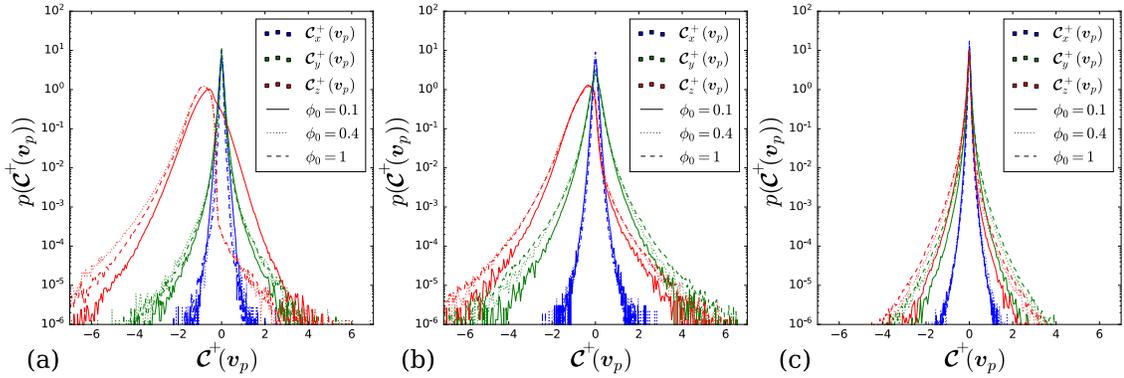
FIGURE 3.17: PDF of three components of curl, normalised by the viscous time scale, for three different mass loadings in the (a) viscous sublayer, (b) buffer layer, and (c) log layer. The line colour indicates the component of the curl, and the line style indicates the Stokes number.

interpretation is not satisfying because the fluid helicity is much less, indicating that the particles are behaving somewhat independently of the fluid in this case. A similar observation was made about the large St^+ particles, and there is also a similarity in the PDF of $\mathcal{C}_z^+$, which is significantly more symmetric at $\phi_0 = 0.1$ than large mass loadings. However, the particle distribution for $\phi_0 = 0.1$ is not random, but organised into streaky structures, suggesting a different cause than particles becoming ballistic and independent of the fluid, as happens at large St^+. This interpretation is confirmed by examining the visualisation of particles coloured by helicity in the viscous sublayer in figure 3.19. At $\phi_0 = 0.1$, the distribution of helicity is coherent, unlike the random speckles of helicity seen in figure 3.9(d) for $St^+ = 60$. The fact that figure 3.19(a,b) are both mostly green (zero helicity) underscores that particle swirling motion is weak regardless of mass loading, and that the sharp difference in PDFs in figure 3.18(a) corresponds to slightly more intense, coherent helicity at $\phi_0 = 0.1$.

This transition warrants further examination, particularly with cases at even smaller mass loading, where the two-way coupling effect becomes negligible. A possible explanation is that the reduced two-way-coupling at small ϕ_0 allows for the transfer of more intensely swirling particle clusters from the buffer layer into the viscous sublayer. This is suggested by the fact that the helicity distribution in the viscous sublayer for ϕ_0 is quantitatively much closer to the distributions in the buffer layer than those in the viscous sublayer at other mass loadings.

3.3.3 Fourier spectra of number density fluctuations

The preceding observations on divergence, curl, and helicity have shown how these quantities are attenuated or enhanced in different flow regions. However, in the PDFs presented, scale information has been lumped together. In this section, spectra of particle number density fluctuations are
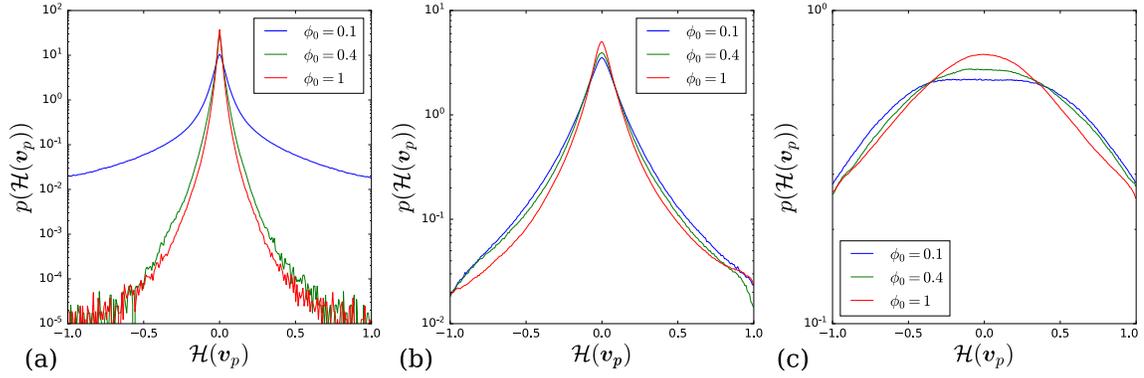
FIGURE 3.18: PDFs of the relative helicity $\mathcal{H}(\boldsymbol{v}_p)$ of the particles for different mass loadings in the (a) viscous sublayer, (b) buffer layer and (c) log layer. Note that the y-ranges differ by orders of magnitude.
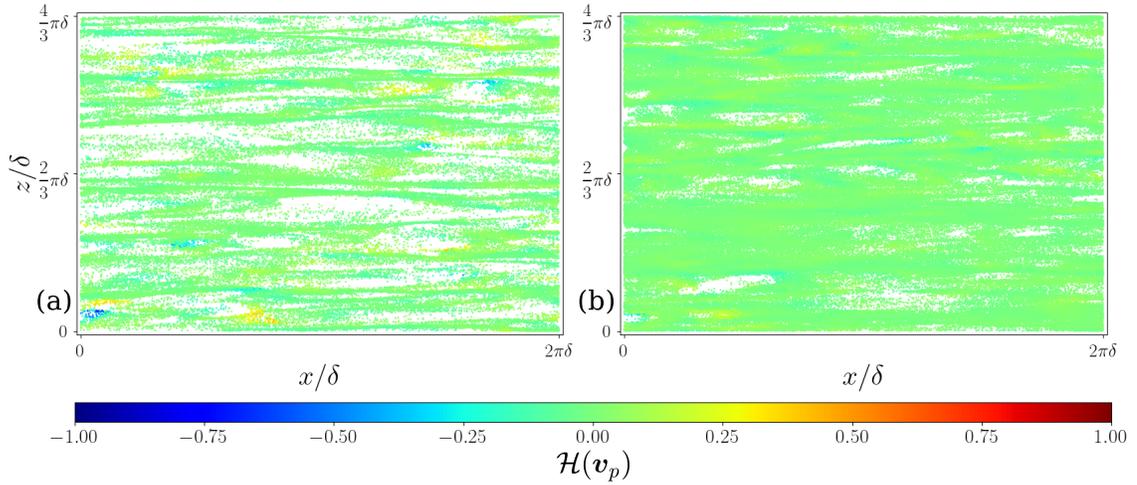
FIGURE 3.19: Particle positions coloured by relative helicity in wall-parallel slices, for (a) viscous sublayer, $\phi_0 = 0.1$; (d) viscous sublayer, $\phi_0 = 0.4$.
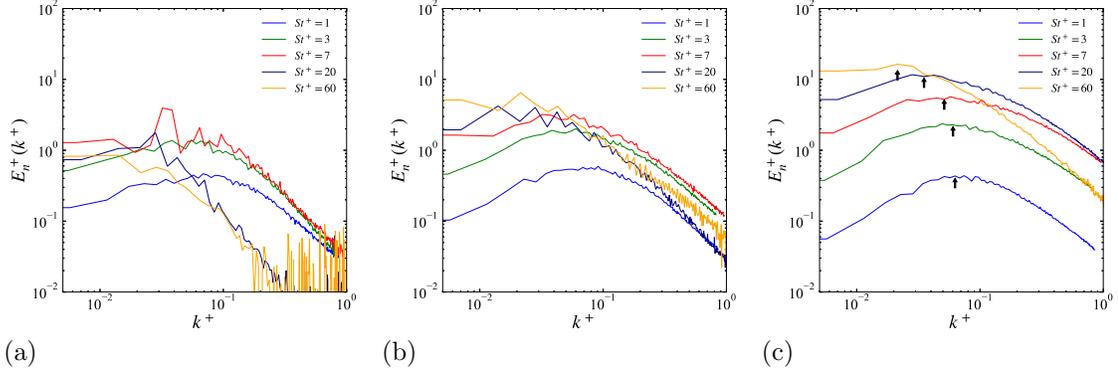
FIGURE 3.20: Fourier spectra of particle number density on 2D slices in (a) viscous sublayer, (b) buffer, and (c) log layers for various Stokes numbers. The spectra and the wavenumber are normalised by the viscous length scale. In (c), the peak locations obtained with a five-point moving average are indicated with arrows.

examined to ascertain which scales of clustering may contribute to the changes observed so far. As in the other sections, the domain is decomposed into wall-parallel slices to characterise the effect of flow region. The location and thickness of the slices are the same as the visualisation in Figure 3.14, i.e. $2 \le y^+ < 3$, $9.5 \le y^+ < 14.5$ and $97.5 \le y^+ < 102.5$ for viscous sublayer, buffer layer, and log layer, respectively. Each spectrum is averaged over 10 snapshots.

The Fourier spectra of number density fluctuations for simulations at different St^+ are shown in figure 3.20. The spectra are normalised by the viscous length scale, i.e., $E_n^+(k^+) = E_n(k^+/l_{visc})/l_{visc}$ and $k^+ = kl_{visc}$.   In the log layer, the Stokes number dependence is similar to that for HIT. Initially as St^+ is increased, the spectra shift vertically while the peak location remains at $k^+ \approx 0.08$, which corresponds to $k\eta \approx 0.2$ (η is the Kolmogorov scale), i.e. slightly smaller than Kolmogorov-scale eddies, which dominate clustering behaviour for small Stokes number. Then for sufficiently large Stokes number ($St^+ > 7$), the spectra move horizontally to smaller wavenumbers. The increased prevalence of large-scale clusters in the log layer for $St^+ = 60$ can be clearly seen by comparing the visualisations in figure 3.7(c,f). This shift is expected at $St_\eta = 1$ per Matsuda et al. [127], and in the present data, it is observed around $St_\eta^{\log} \approx 0.5$. A comparison of peak location vs. Stokes number normalised by Kolmogorov scale is shown in figure 3.21 for both the present log layer data and the HIT data of Matsuda et al. [127]. Given the differences in Reynolds number, two-way coupling, and flow anisotropy, we find that the quantitative similarity is striking. As the peak location begins to shift significantly, the slope of the spectra at high wavenumbers in the log layer also decreases in qualitative agreement with the results from [127]. This change in slope corresponds to the reduction in fine-scale clusters observed going from $St^+ = 3$ to $St^+ = 60$ in figure 3.7(c,f).

In the buffer layer, a similar horizontal shift to larger scales as St^+ is increased can also be observed, again corresponding to the increased prevalence of large-scale clusters and reduction in
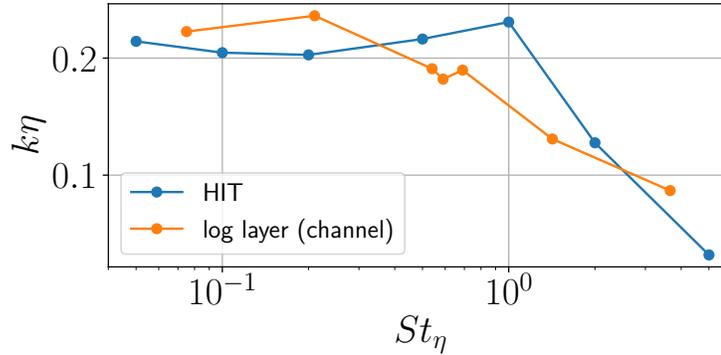
FIGURE 3.21: Wavenumber of maximum in spectra of particle number density fluctuations as a function of Kolmogorov Stokes number for HIT [127] and the logarithmic layer in the present dataset.

small-scale clusters seen in figure 3.20(b,e). In the viscous sublayer, the spectra are noisier, particularly at large St^+, because the particles become more randomly distributed, and the periodogram, used to estimate the spectrum, is an inconsistent estimator due to the presence of oscillations [166].

Figure 3.22 shows the spectra as a function of ϕ_0. In general, there is an attenuation at all scales across all layers as mass loading is increased. This is expected from the turbulence attenuation caused by two-way coupling (figure 3.2(c,d)). Similar to the PDFs of divergence, curl, and helicity, the attenuation in spectra is most dramatic in the viscous sublayer, except at the very largest scales (small wavenumbers). This suggests that the attenuation of clustering, vortical motion, and swirling as mass loading increases originate from small to moderate scales. The dominant scales in the viscous sublayer correspond to the length of and spacing between the low-speed streaks, which do not change drastically as ϕ_0 is increased from 0.1 to 1.0. In the buffer and log layers, the peak location does not shift much, which is expected because St^+ is approximately constant for the mass loading variation. This is confirmed by the visualisations in figure 3.14, which show similar scales of clustering, albeit with greater number density of particles, as ϕ_0 is increased.

3.4 Conclusions

In this study, tessellation-based methods for computing particle velocity divergence, curl, and helicity were applied to turbulent channel flow, and the influence of mass loading and Stokes number in various flow regions was examined in detail.

Probability density functions (PDFs) of particle velocity divergence and curl reveal that these quantities vary most intensely in the buffer layer. Since turbulent kinetic energy peaks in the buffer layer, this implies that strong velocity fluctuations drive more intense cluster formation/destruction, as well as vorticity. In the viscous sublayer, particles form elongated clusters in the low speed steaks,
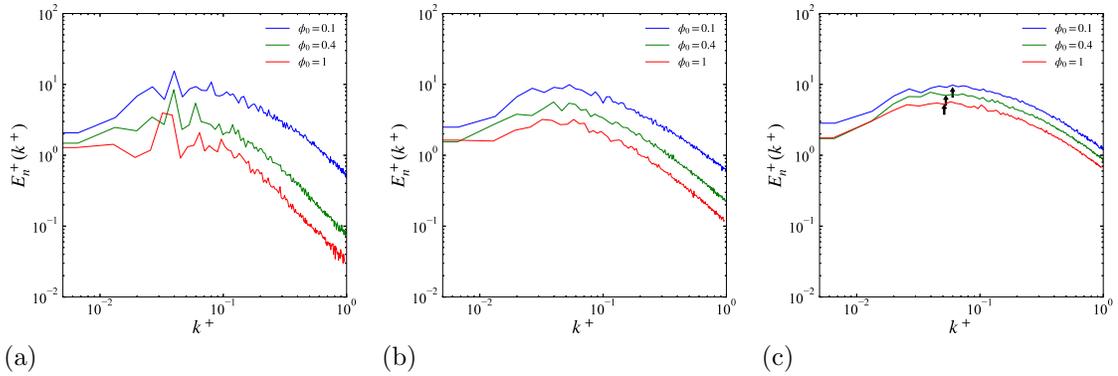
FIGURE 3.22: Fourier spectra of particle number density on 2D slices in (a) viscous sublayer, (b) buffer, and (c) log layers for various mass loadings. The spectra and the wavenumber are normalised by the viscous length scale. In (c), the peak locations obtained with a five-point moving average are indicated with arrows.

which are transported without converging or diverging. The wall boundary condition naturally results in a flow with less swirling motion than in the rest of the channel, but the presence of particles stabilises the viscous sublayer and makes the streaks longer-lived. This makes both the fluid and particle motions even more two-dimensional, except in the case of very large Stokes number.

In the logarithmic layer, many similarities with results for homogeneous isotropic turbulence (HIT) reported by Oujia et al. [151] were observed. Particle clustering patterns, as well as distributions of divergence, curl, and helicity are qualitatively similar to one-way coupled particles in HIT, despite significant force coupling from the particle phase, the turbulence anisotropy that results from two-way coupling in wall-bounded flows, and quantitative differences in Reynolds number. Further, the spectra of particle number density fluctuations exhibit a similar dependence on Stokes number as compared to the HIT dataset of Matsuda et al. [127].

The dynamics of particle clusters in channel flow are affected similarly by Stokes number as they are in HIT, except in the viscous sublayer, where the relatively low turbulence intensity and proximity to wall boundary result in particles which behave more independently of the fluid. Like in HIT, increasing the Stokes number results in more intense convergence/divergence, as well as more intense vortical motion. This effect eventually saturates, except in the viscous sublayer, where particles become more independent of the relatively quiescent fluid flow. Due to two-way coupling, the particle motion becomes more two-dimensional (zero-helicity) in the log and buffer layers as Stokes number is increased. In the viscous sublayer, an initial attenuation is followed by an increase in swirling motion as St^+ is increased. This non-monotonic trend further illustrates the stark differences in clustering between the viscous sublayer and other flow regions.

The effect of the mass loadings studied in this work is more subtle than the effect of Stokes number, but again, the viscous sublayer shows unique cluster dynamics. The changes observed in the

viscous sublayer as mass loading is increased are that the variance of particle velocity divergence and curl are both decreased, and that swirling particle motions are strongly attenuated. These changes are expected from two-way coupling with the fluid, which results in more anisotropic turbulence and more persistent low-speed streaks near the wall. It is surprising that similar effects are not observed in the buffer and log layers, given that the inertia of fluid and particle phases are comparable at $\phi_0 = 1$, and the strong attenuation of the fluid-phase Reynolds stress that occurs over this range (figure 3.2(d)). Nevertheless, as mass loading increases from $\phi_0 = 0.1$ to $\phi_0 = 1.0$, there are only small changes to divergence, curl, and helicity in the buffer and log layers. Likewise, the spectra of number density fluctuations are much more strongly affected in the viscous sublayer as mass loading is increased than in other parts of the flow. We also acknowledge the possibility that further increasing the mass loading above $\phi_0 = 1$ would lead to greater changes in cluster dynamics throughout the flow.

The PDFs based on particle motions not only quantify the cluster formation/annihilation, rotation, and swirling of particle clouds, but also are consistent with the structures seen in the visualisations, as well as the number density spectra. Thus, the present approach provides both statistical and structural insights into particle clustering in wall turbulence, which mean profiles and Eulerian analysis cannot do alone. The techniques used in this paper can also be applied to study clustering, rotation, and swirl of particle clouds in other regimes in the vast parameter space of particle-laden flows [31]. For example, in flows with larger particle volume fraction, collisions are expected to play a significant role, and their influence on the clustering, rotation, and swirling of particle clouds should be examined. In future work, wavelet-based statistics and multi-resolution analysis of particle clustering [128] can be used in conjunction with these techniques to better characterise the scale dependence of particle clustering, and how different scales of clusters contribute to the dynamics of particle clouds.

3.5   Appendix A: Visualizations of fluid velocity

To gain a better sense of the changes in flow structure, instantaneous snapshots of streamwise velocity in the viscous sublayer, buffer layer and log layer are shown in figure 3.23, which compares two values of Stokes number St^+, and figure 3.24, which compares two values of mass loading ϕ_0. These show the same time snapshots as figures 3.7 and 3.14, respectively. The effect of increasing ϕ_0 from 0.1 to 0.4 is subtle, but in general results in an attenuation of smaller scales in the fluid velocity field. In the viscous sublayer and buffer layer, this manifests as more persistent and slightly more widely spaced near-wall streaks. The persistence and stability of streaky structures illustrates how particles stabilise the near-wall flow and dampen velocity fluctuations. The effect of changing St^+ from 3 to 60 is more extreme, but qualitatively similar. Near-wall streaks become much more stable and wider at large St^+, and the streaks persist well into the log layer, which indicates that

large St^+ particles can act to thicken the viscous sublayer, make the flow more two-dimensional, and eventually kill the turbulence. The thickening of the viscous sublayer at large St^+ can also be seen in the mean velocity profile in figure 3.2(a).
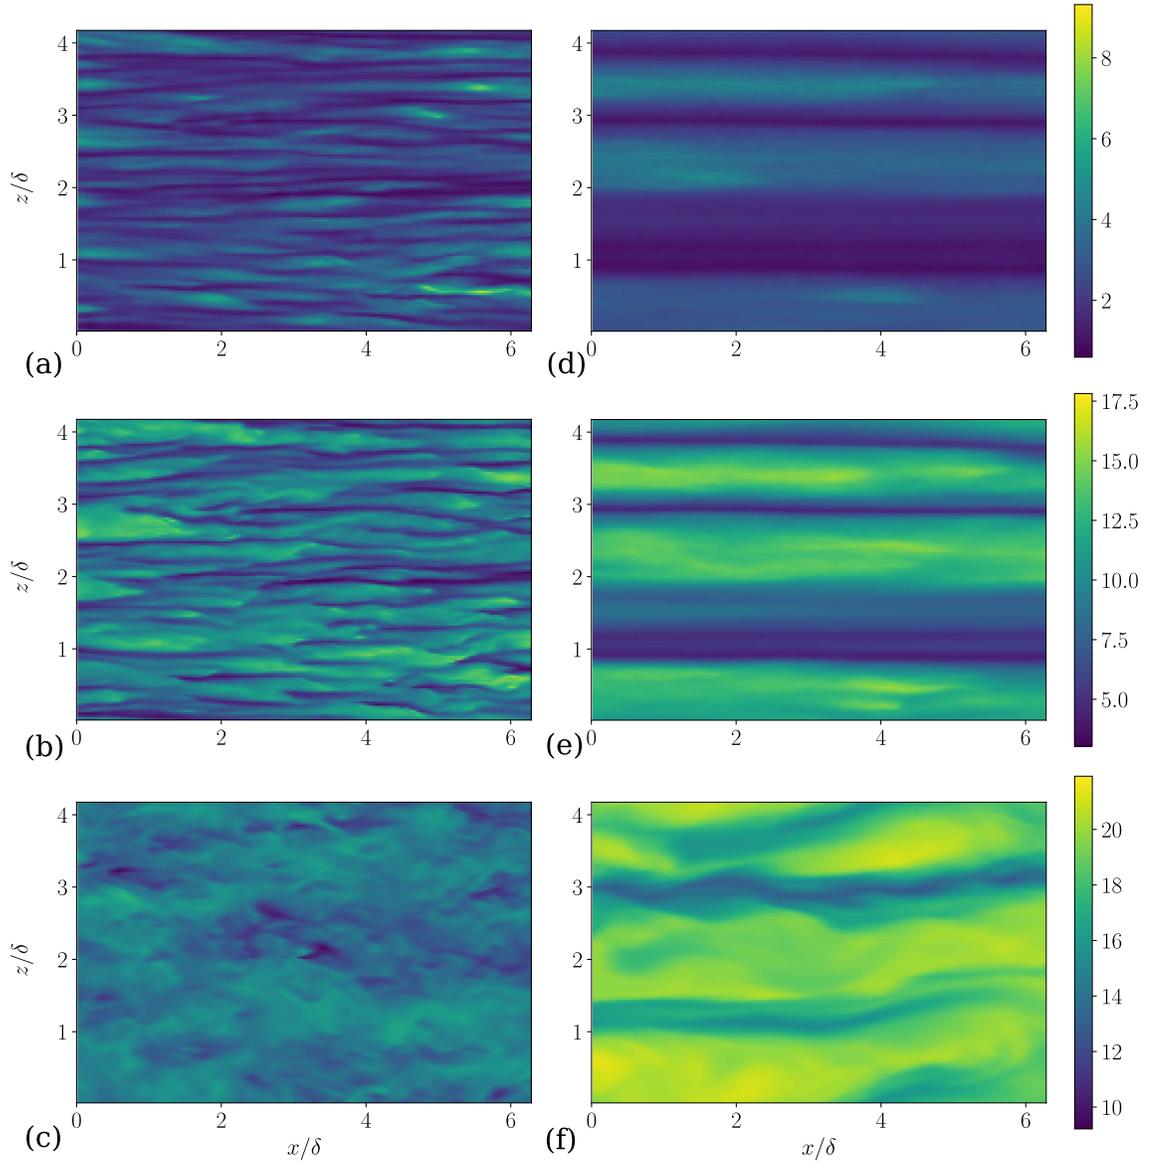
FIGURE 3.23: Fluid streamwise velocity u^+ for (a) $\phi_0 = 1.0$, $St^+ = 3$ at $y^+ = 2.5$; (b) $\phi_0 = 1.0$, $St^+ = 3$ at $y^+ = 12$; (c) $\phi_0 = 1.0$, $St^+ = 3$ at $y^+ = 100$; (d) $\phi_0 = 1.0$, $St^+ = 60$ at $y^+ = 2.5$; (e) $\phi_0 = 1.0$, $St^+ = 60$ at $y^+ = 12$; (f) $\phi_0 = 1.0$, $St^+ = 60$ at $y^+ = 100$.
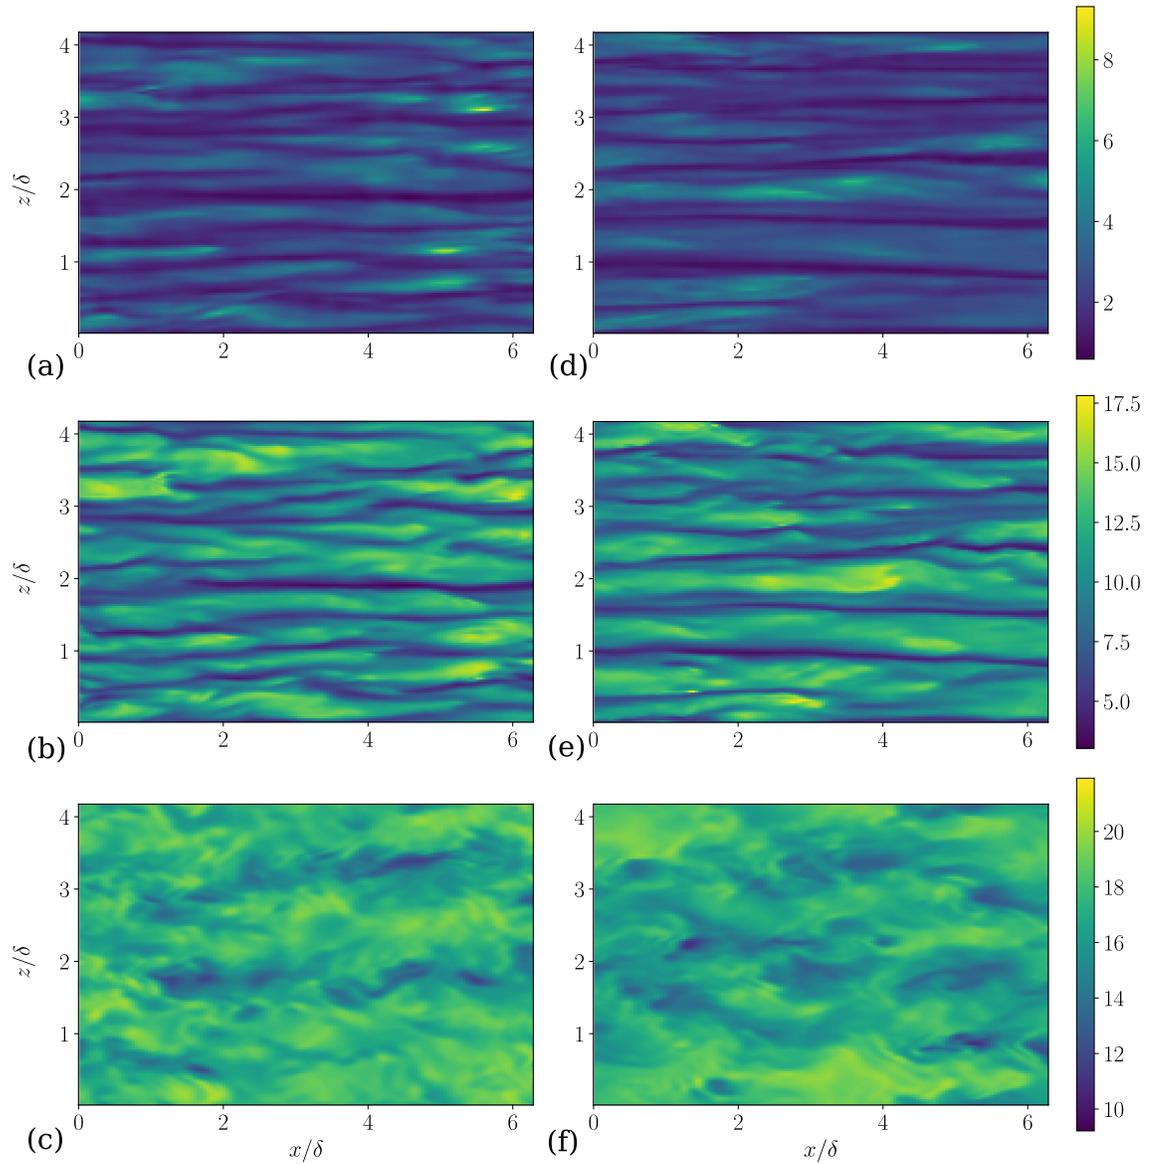
FIGURE 3.24: Fluid streamwise velocity u^+ for (a) $\phi_0 = 0.1$, $St^+ = 7$ at $y^+ = 2.5$; (b) $\phi_0 = 0.1$, $St^+ = 7$ at $y^+ = 12$; (c) $\phi_0 = 0.1$, $St^+ = 7$ at $y^+ = 100$; (d) $\phi_0 = 0.4$, $St^+ = 7$ at $y^+ = 2.5$; (e) $\phi_0 = 0.4$, $St^+ = 7$ at $y^+ = 12$; (f) $\phi_0 = 0.4$, $St^+ = 7$ at $y^+ = 100$.

Chapter 4

Radiatively heated particle-laden channel flow

4.1 Introduction

Particle-laden flows are widespread in both natural and industrial applications, including such diverse flows as volcanic plumes, sediment transport in rivers, sandstorms, chemical process reactors, and exhaust scrubbing at coal-fired power plants. Some of these examples may include significant temperature variation in both the fluid and particle phases. In these settings, there can be many complex interactions among heat transfer, turbulence, and particles. Another system in which this type of coupling is expected to be important is a novel device for harnessing solar energy, originally proposed by Hunt [90] and Abdelrahman et al. [1], called a small-particle solar receiver. This device aims to overcome limitations associated with existing working fluids used in concentrating solar power plants [85]. Variants of this device have been proposed, including with falling particles in an open cavity, stationary particles in a transparent tube, and fluidized particles in an opaque tube [84]. If solid particles are used and an appreciable amount of radiation is absorbed by them, then it is likely that a small-particle solar receiver would need to operate in a regime with large mass loading, which implies strong momentum coupling between the phases. Clearly, there will also be strong heat transfer coupling, and turbulence will give rise to complex interactions among these physics. For this reason, we present here a study of radiatively heated channel flow at large mass loading and radiation intensity, to better understand the sorts of complex couplings that might arise in such a device, and what implications these might have for design.

Previous studies of radiation-particle-turbulence interaction have tended to examine cases with low mass loading and radiation intensity, lower than would be expected in a practical device. The

effect of preferential concentration on absorption was studied [164] and found to diminish the absorption efficiency. This is due to decreased optical depth as a consequence of particle clustering [64]. A number of studies have examined particle-radiation-turbulence interaction in homogeneous isotropic flows. Pouransari and Mani [165] studied the dimensionless groups affecting particle-turbulence-heat transfer coupling in this setting and found that Stokes number and another parameter related to heat transfer in the gas were the dominant parameters. Additional effects arise due to interaction with gravity, both through altered particle settling [66] and turbulence induced when particles heat a surrounding buoyant fluid [215]. A benchmark experimental and simulation comparison of radiatively heated particle-laden duct flow was also conducted by Esmaily et al. [55].

The present study aims to push the envelope of simulations toward realistic mass loadings and radiation intensities, while maintaining a simplified channel flow geometry, modest Reynolds number, and low volume fraction to control computational cost. The remainder of this Chapter is organized as follows. Section 4.2 describes the computational methods and simulation cases used. Section 4.3 reports on the effects of heating on various flow statistics. Section 4.4 describes and evaluates a simple 1D model, and section 4.5 provides a summary and conclusion.

4.2 Methods

The Soleil-MPI code is used to solve the coupled equations for fluid flow, particle motion, and radiation heat transfer [55]. The numerical methods in the heated flow are largely the same as for the isothermal domain, discussed in Chapter 2. In this section we highlight the modifications to the solver in order to simulate heated problems.

The treatment of radiation heat transfer in the Soleil-MPI code is presented in detail in Frankel [65]. The radiation transport equation (RTE) is solved in quasi-steady form, i.e. the steady state solution based on the current state of the particle-laden flow. This is done because the timescale of radiation heat transfer (the speed of light) is much faster than any flow or particle timescale. Furthermore, we assume a monochromatic radiation source and consider only purely absorbing particles, neglecting both scattering [1] and blackbody emission [2]. These assumptions significantly simplify the RTE, as there is no angular variation in radiation intensity, nor spectral variation

---

[1] Scattering is neglected because it does not affect the Eulerian evolution equation for the thermal energy of particle clouds (see Frankel [65], equation 2.37). However, from the standpoint of an individual particle, it is expected to be relevant as the values of absorption and scattering coefficients calculated by Frankel [65] in section 5.2.3 are $Q_a = 0.4$ and $Q_s = 0.7$, respectively.

[2] The relative importance of blackbody radiation can be estimated by considering the balance of blackbody radiation and absorption on a single particle. Assuming a spherical particle and a $0\,[K]$ ambient, the ratio of blackbody to absorption is $4\sigma T_p^4/I_0$, where σ is the Stefan-Boltzmann constant, T_p is the particle temperature, and I_0 is the radiation intensity at the boundary. In the simulation case with the greatest particle heating ($\phi_0 = 2$), the maximum particle temperature occurs in a thin layer near the wall, for which which $T_p \approx 1100\,[K]$, for which blackbody re-radiation is 21% the intensity of absorbed radiation. The maximum particle particle temperature in the centerline is $T_p \approx 700$, for which the ratio if 3.5%. Based on this analysis, it is expected that blackbody radiation will be insignificant in the channel centerline across all simulation cases, and will play a significant, but not dominant role very close to the walls at large mass loading. Thus, we neglect blackbody radiation for simplicity

in radiation intensity (i.e. gray radiation). In other words, $I = I(x, y, z)$, where I is the radiation intensity. With all these simplifications, the RTE for a collimated source aligned with the y direction (i.e. wall-normal) can be written as

$$\frac{dI}{dz} = -\sigma_a I \tag{4.1}$$

$$\sigma_a = Q_{abs} \frac{\pi D_p^2}{4} n \tag{4.2}$$

where Q_{abs} is the absorption coefficient of the particles. We have assumed that the carrier fluid is transparent to radiation, as is appropriate for solar radiation in air. While this simplified form can be solved exactly for a given field $n(x, y, z)$, Soleil-MPI is designed for more general cases, and so employs the discrete ordinates method [62] to solve it numerically. Since the number density n is a continuum property, and the particle positions are discrete, the particle positions are represented as n on the fluid grid using the homogenization procedure described in Frankel [65] and Frankel et al. [64] in order to reduce numerical errors.

The radiation intensity $I(x, y, z)$ obtained by solving equation 4.1 couples to the other physics solvers (fluid momentum and particle tracking) via radiation absorbed by the particle phase. Modifications to the particle-tracking ODEs as a consequence of radiation heat transfer are described next. In addition to the ODEs for particle velocity and position (equations 2.4 and 2.5), an additional ODE is solved for energy conservation of each particle in the heated domain:

$$m_p C_{vp} \dot{T}_p^{(m)} = Q_{abs} \frac{\pi}{4} D_p^2 I - e^{(m)} \tag{4.3}$$

where C_{vp} is the (constant) specific heat of each particle, $T_p^{(m)}$ is the temperature of particle m, and $e^{(m)}$ is the convective heat exchange term with the fluid. The term involving I is the radiation absorption source term. The convective heat exchange term is given by

$$e^{(m)} = \pi D_p^2 h_p \left(T_p^{(m)} - \widetilde{T}^{(m)} \right) \tag{4.4}$$

where $\widetilde{T}^{(m)}$ is the undisturbed fluid temperature at the location of particle m, and h_p is the convective heat transfer coefficient, computed according to $Nu = \frac{h_p D_p}{k} = 2$, where Nu is the Nusselt number for a sphere, and k is the fluid thermal conductivity [3]. The temperature $\widetilde{T}^{(m)}$ is computed with tri-linear interpolation, without a two-way coupling correction [4].

Finally, the Navier-Stokes solver for fluid mass and momentum is also affected by particle heating. While the heat transfer is significant in these flows, they are still low-speed, and so we employ the low-Mach approximation [121]. Under this assumption, the pressure is separated into thermodynamic

---

[3] The assumption that $Nu = 2$, used here for simplicity, is valid for small particle Reynolds numbers and small heat differences between fluid and particle. Ganguli and Lele [68] recently relaxed these assumptions and obtained improved correlations for Nu and drag coefficient C_D.

[4] Since the simulations presented in this work were conducted, correction procedures for the undisturbed fluid temperature, analogous to those for momentum, have been developed by Liu et al. [119] and Horwitz et al. [88].

and hydrodynamic components: $p = p_0 + \tilde{p}$. The thermodynamic pressure p_0 is assumed to be constant and to obey the ideal gas law

$$p_0 = R_g \rho T, \tag{4.5}$$

where R_g is the specific gas constant for air, and T is the fluid temperature. The low-Mach, variable density versions of the equations for mass and momentum are given by

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \boldsymbol{u}) = 0 \tag{4.6}$$

$$\frac{\partial (\rho \boldsymbol{u})}{\partial t} + \nabla \cdot (\rho \boldsymbol{u} \otimes \boldsymbol{u}) = -\nabla \tilde{p} + \nabla \cdot \left(\mu \left(\nabla \boldsymbol{u} + \nabla \boldsymbol{u}^T - \frac{2}{3} \left(\nabla \cdot \boldsymbol{u} \right) \boldsymbol{1} \right) \right) + \sum_{m=1}^{N_p} \boldsymbol{f}^{(m)} \delta_{\text{Dirac}}(\boldsymbol{x} - \boldsymbol{x}_p{}^{(m)}),$$
$$\tag{4.7}$$

where $\boldsymbol{1}$ is the identity tensor. These equations are solved using the spatial derivative and explicit time-stepping schemes described for the momentum equation in Chapter 2. This system of equations is closed by a Poisson equation for the hydrodynamic pressure $\tilde{p}$. The Poisson equation is obtained by considering the low-Mach transport equation for fluid thermal energy

$$\frac{\partial (\rho C_v T)}{\partial t} + \nabla \cdot (\rho C_p T \boldsymbol{u}) = \nabla \cdot (k \nabla T) + \sum_{m=1}^{N_p} e^{(m)} \delta_{\text{Dirac}}(\boldsymbol{x} - \boldsymbol{x}_p{}^{(m)}), \tag{4.8}$$

where C_v and C_p are the constant-volume and constant-pressure specific heats of the fluid, respectively. Given the low-Mach assumption, ρT is a constant, and the above equation can be reduced to an elliptic equation relating fluid velocity and temperature

$$\nabla \cdot \boldsymbol{u} = \frac{R_g}{C_p p_0} \left\{ \nabla \cdot (k \nabla T) + \sum_{m=1}^{N_p} e^{(m)} \delta_{\text{Dirac}}(\boldsymbol{x} - \boldsymbol{x}_p{}^{(m)}) \right\}. \tag{4.9}$$

Instead of a divergence-free velocity field, the divergence depends on heat conduction and heat transfer from the particles. The Poisson equation is solved using the same methods described in Chapter 2, but the heating source term from the particles changes the right hand side.

Throughout the preceding equations, the fluid thermal conductivity k and dynamic viscosity μ are not assumed to be constant, and they depend on temperature. The thermal conductivity is computed according to the correlation for air from Rumble [170]

$$k = aT^3 + bT^2 + cT + d, \tag{4.10}$$

where $a = 1.5207 \times 10^{-11}$ $[Wm^{-1}K^{-4}]$, $b = -4.8574 \times 10^{-8}$ $[Wm^{-1}K^{-3}]$, $c = 1.0184 \times 10^{-4}$ $[Wm^{-1}K^{-2}]$, $d = -3.9333 \times 10^{-4}$ $[Wm^{-1}K^{-1}]$. The dynamic viscosity is computed according to Sutherland's
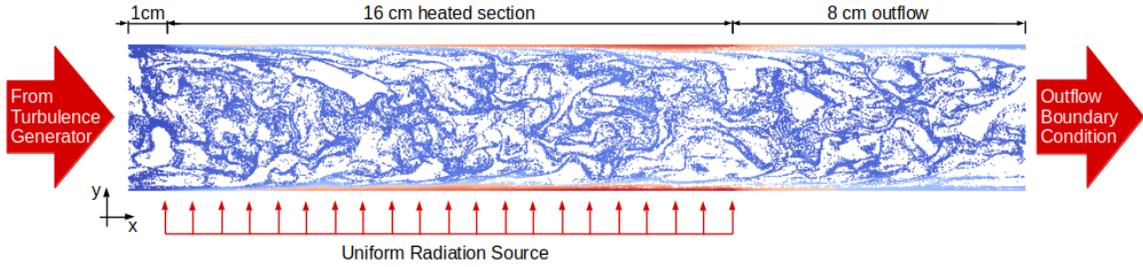
FIGURE 4.1: Schematic of heated simulation domain, shown in the $x - y$ plane.

formula [191]

$$\mu = C_1 \frac{T^{3/2}}{T + S},\tag{4.11}$$

where $C_1 = 1.51065 \times 10^{-6} \ [Pa \ sK^{-0.5}]$ and $S = 120 \ [K]$.

While the dimensions of the heated channel are the same as the unheated channel, the radiation heat transfer causes streamwise development in the flow, and so inflow and outflow boundary conditions must be used in the streamwise (x) direction. The no-slip wall boundary conditions at $y = [0, 2\delta]$ and periodic boundaries in z are the same as the unheated domain. The inflow boundary condition to the heated domain is provided by a synchronized unheated simulation, described in Chapter 2. A y-z slice of the unheated fluid domain is provided at each timestep as the inflow for the fluid, and the particles are advected into the heated domain according to their velocity at the periodic boundary in the unheated domain. An outflow boundary condition for each flow variable (density and velocity) is applied such that a convective velocity equal to the bulk velocity is imposed. The radiation boundary condition is applied in the range $0.5 \leq x/\delta \leq 8.5$, as illustrated in figure 4.1. This long outflow length ensures the boundary condition does not contaminate the solution in the heated region. The boundary condition for the fluid density at the walls is adiabatic. In other words, while the walls are transparent to radiation, there are no conductive or convective heat losses to the environment in this idealized setup.

The same domain dimensions and grid are used for the heated and unheated domains: $[Nx, Ny, Nz] = [280, 140, 140]$. Because the viscosity increase due to heating, the Reynolds number is lower in the heated domain, and the flow is thus better-resolved. The same time step is used for both heated and unheated domains, and the timestep is determined based on a CFL condition. As a result, the timestep required is smaller for the larger values of mass loading, which result in more heat absorption and viscosity increase. In these flows, to reduce computational cost, collisions are not computed at every timestep. However, the physical time between collision computations is still always less than the timestep required for stability in the unheated domain.

The parameters of the unheated domain (also known as the turbulence generator) are shown in table 4.1. These are the same as Flows 1-4 in Chapter 2. The parameters specific to the heated

	Re_τ	N_p	ρ_p/ρ	ϕ_0	St^+
Flow 1	225	1.43×10^7	7.42×10^3	0.1	6.90
Flow 2	227	5.70×10^7	7.42×10^3	0.4	7.01
Flow 3	234	1.43×10^8	7.42×10^3	1.0	7.45
Flow 4	258	2.85×10^8	7.42×10^3	2.0	9.04

TABLE 4.1: Parameters of unheated inflow associated with the heated channel flow cases: the friction Reynolds number Re_τ, the number of particles N_p, the particle to fluid density ratio ρ_p/ρ, the overall mass loading ϕ_0, and the friction Stokes number St^+.

$I_0\,[MW/m^2]$	$C_v\,[J/kg\text{-}K]$	$C_p\,[J/kg\text{-}K]$	$p_0\,[MPa]$	$C_{vp}\,[J/kg\text{-}K]$	Q_{abs}
1.5625	732	1012	0.1	450	0.4

TABLE 4.2: Simulation parameters specific to the heated simulations: the radiation intensity I_0, the constant-volume specific heat of the fluid C_v, the constant-pressure specific heat of the fluid C_p, the hydrodynamic pressure p_0, the constant-volume specific heat of the particles C_{vp}, and the particle radiation absorption coefficient Q_{abs}.

domain simulations are given in table 4.2. These parameters are shared by all heated flow cases, and are based on the experiments using nickel particles in air described in Esmaily et al. [55]. Statistics are obtained by running the simulation until steady state is reached, and then averaging over a minimum of five flow-through-times, as well as in the spanwise direction (z).

4.3   Results

In this section, for illustration purposes, we focus on Flow 3 ($\phi_0 = 1$), because the flow changes are qualitatively similar across the chases, differing primarily in the quantity of heat absorption. After applying the radiation boundary condition, the particles absorb radiation and heat up. Then, they transfer energy to the cooler carrier gas via convective exchange. Because of the constant thermodynamic pressure, the heating of the gas translates directly into density decrease (i.e expansion). Instantaneous snapshots of the fluid density in the $x - z$ plane are plotted in figure 4.2, to illustrate differences in streamwise development near the wall versus the centerline. In the buffer layer, the density decreases more sharply than in the centerline. Further, the structure of the density field in the buffer layer is dominated by streamwise streaks, which originate from the low-speed streaks in the fluid phase. In particle-laden flows, particles tend to accumulate within the low-speed streaks, and if the flows are two-way-coupled, the particles influence the persistence and spacing of the streaks, as illustrated in figures 3.23 and 3.24. Since the particles are the heat source for the fluid, a locally high concentration of particles results in a locally intense heat source. Thus, the low-speed streaks from the unheated domain correspond to the lower-density streaks in the heated domain. In the centerline, the fluid expansion is less intense and more gradual. Further, the structure of the density
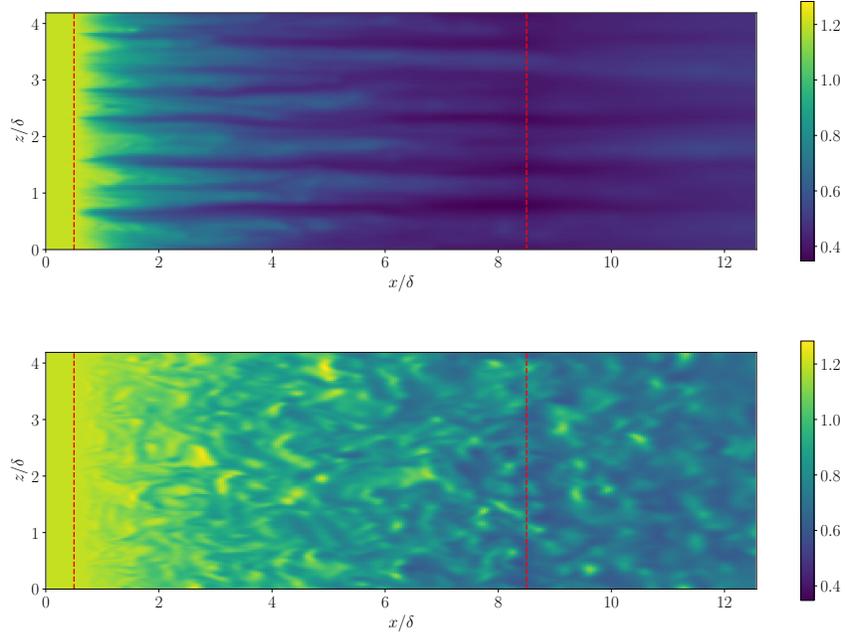
FIGURE 4.2: Snapshots in $x-z$ planes of fluid density (dimensions kg/m^3) for the case with $\phi_0 = 1$, in the buffer layer (top panel, $y_0^+ = 12$) and the channel centerline (bottom). The beginning and end of the heated section are each marked with a dashed red line.

field lacks directional dependence. This is due to the more isotropic distribution of particles (heat sources), like those in the logarithmic layer shown in figure 3.14(c,f).

Because the flow is confined in a channel, the fluid expansion naturally results in flow acceleration, as illustrated in the fluid velocity plotted in figure 4.3. Like the fluid density, the fluid velocity is dominated by streaky structures in the buffer layer and more isotropic structures in the centerline. By comparing figures 4.2 and 4.3, it is clear that the slower streaks in the fluid correspond to the lower density and also that the contrast between lower and higher speeds across the z-axis decreases with streamwise development. This is because the low speed streaks carry more particles and therefore absorb more thermal energy. Thus, the fluid in the low-speed streaks expands and accelerates more than the surrounding fluid, resulting in the buffer layer becoming more homogeneous in the spanwise direction. On top of this broad homogenization, it appears that some streaks may merge or vanish with streamwise development. This is most clearly seen in the buffer-layer velocity (figure 4.3), where the streak beginning at $y/\delta \approx 1$ fades away at $x/\delta \approx 5$, or merges into the streak beginning at $y/\delta \approx 0.7$. Similarly, the streak beginning at $y/\delta \approx 3.3$ seems to simply vanish. In the centerline the flow acceleration is substantial, and the change in velocity due to acceleration is a more significant source of velocity variation than eddies, which are the cause of spanwise inhomogeneity. These changes illustrate how strong radiation heating can make the flow more uniform and less influenced
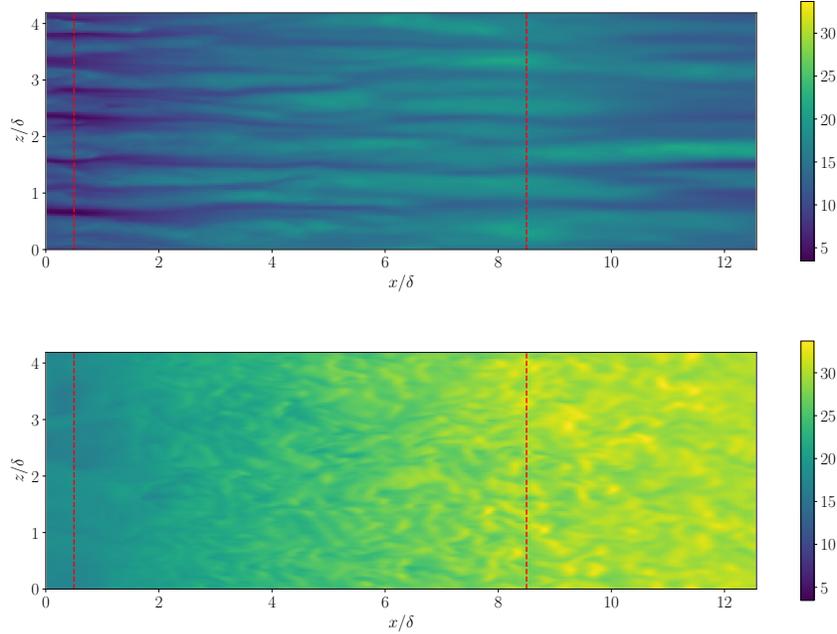
FIGURE 4.3: Snapshots in $x - z$ planes of fluid streamwise velocity, normalized by the viscous scale from the unheated flow for the case with $\phi_0 = 1$ in the buffer layer (top panel, $y_0^+ = 12$) and the channel centerline (bottom). The beginning and end of the heated section are each marked with a dashed red line.

by turbulence, a fact which will be further illustrated when we turn to velocity fluctuations in section 4.3.3.

As the fluid expands and accelerates, so do the particle clusters transported by the fluid, due to the momentum coupling between the phases. As a result, the particle number density decreases with streamwise development, and more sharply so near the wall than in the centerline. To illustrate how the phases develop differently in the streamwise direction, we plot snapshots of the local mass loading in figure 4.4. As expected, the mass loading is higher near the wall and dominated by streaky structures, while in the centerline the particle structures are more isotropic. Additionally, the regions of higher mass loading correspond to the regions of larger fluid expansion, because higher mass loadings correspond to more localized heat sources. What is most striking about the mass loading is how there appears to be little change in the structures when the radiation boundary condition is applied and as the flow develops in the streamwise direction. This indicates that the particles and fluid phases expand together in tight synchronization. It is expected that particle expansion and acceleration should lag behind the fluid because of the particles' inertia, but the Stokes number of these simulations ($St^+ \approx 8$) is low enough that this lag is not evident, even as the fluid density in the buffer layer goes through a sharp transition in the range $x/\delta < 2$. As the Stokes
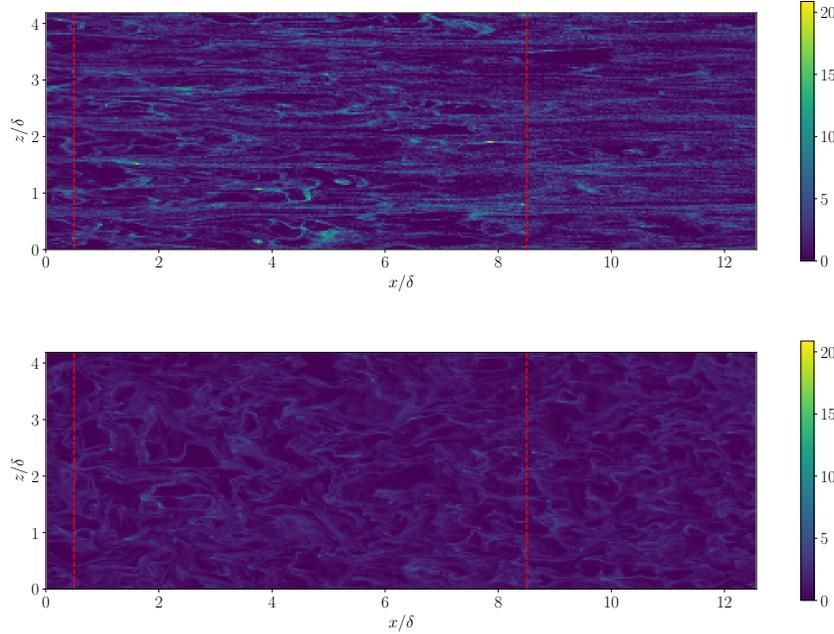
FIGURE 4.4: Snapshots in $x - z$ planes of particle mass loading for the case with $\phi_0 = 1$ in the buffer layer (top panel, $y_0^+ = 12$) and the channel centerline (bottom). The beginning and end of the heated section are each marked with a dashed red line.

number increases, we conjecture that this lag would become evident, at least in the sharp transition region after the radiation boundary condition is applied.

Interestingly, even as there are some signs of turbulence attenuation in the snapshots of fluid properties, the mass loading (and therefore the particle number density) field remains inhomogeneous, exhibiting similar structures due to preferential concentration as the flow accelerates. This suggests that while the eddies which generate these filamentary structures may be weakening, the inhomogeneous distribution may remain frozen as it is advected by a decreasingly turbulent flow.

The changes to the flow discussed so far are further illustrated in Reynolds-averaged (time and spanwise averaged) profiles of fluid density in figure 4.5 and mass loading loading in figure 4.7. In addition to dependence on the wall-normal coordinate, these profiles also show the difference between bottom (where the radiation source is located) and top walls. As radiation travels across the height of the channel, it is attenuated as particles absorb its energy. As a result, the radiation intensity is weaker when it reaches the top wall, and all the effects described so far - heating, expansion, acceleration - are weaker. The amount of difference between the near and far walls is directly related to the amount of radiation absorbed by the particles, i.e. the optical depth of the mixture. In the present suite of simulations, this means that greater differences between the walls are observed as mass loading is increased.

4.3.1 Non-dimensionalization

It is important clarify that in the following discussion, the non-dimensionalization used for the fluid velocity and y-coordinate depends, like the flow itself, on x. The normalization is based on the wall quantities as they develop in the streamwise direction. We make use of the wall quantities ρ_w, μ_w, and τ_w

$$\rho_w(x) = \overline{\rho}|_{y=0} \tag{4.12}$$

$$\mu_w(x) = \overline{\mu}|_{y=0} \tag{4.13}$$

$$\tau_w(x) = \overline{\mu \left(\frac{\partial u}{\partial y} \right)}\Bigg|_{y=0} \tag{4.14}$$

where the overbar denotes Reynolds averaging in homogeneous directions (time and spanwise coordinate here). From these we compute the normalized (u^+, y^+) coordinates as

$$u^+(x,y) = \frac{\overline{u}}{\sqrt{\tau_w/\rho_w}} \tag{4.15}$$

$$y^+(x,y) = \frac{y\sqrt{\rho_w \tau_w}}{\mu_w} \tag{4.16}$$

These have been written for the bottom wall, but the same principle is applied for plotting profiles from the top wall, except that wall quantities are evaluated at $y = 2\delta$.

4.3.2 Mean quantities

Figure 4.5 illustrates the sudden density decrease near the wall observed from the snapshots in figure 4.2. A feature that comes out more clearly in the Reynolds-averaged plots is that the centerline flow behaves oppositely: there is initially no density decrease, and eventually the density change near the centerline accelerates to match the rate of change near the wall, i.e. the curves become more equally spaced along their length. This lag and catch-up suggests that turbulent mixing plays an important role in redistributing the heat which is preferentially absorbed near the wall. Without this mechanism, the near-wall region would continue to heat up at a faster rate than the centerline. Finally, we note that the fluid density profiles differ slightly between top and bottom walls, with the greatest expansion happening at the bottom wall. Because of radiation attenuation, the peak in the density profile (i.e. the coldest fluid) is shifted closer to the top wall than the bottom.

The mean temperature profiles, shown in figure 4.6, show qualitatively the same behavior as the fluid density. This is to be expected from the low-Mach approximation, which results in

$$\overline{T} = \frac{\frac{P_0}{R_g} - \overline{\rho' T'}}{\overline{\rho}}. \tag{4.17}$$
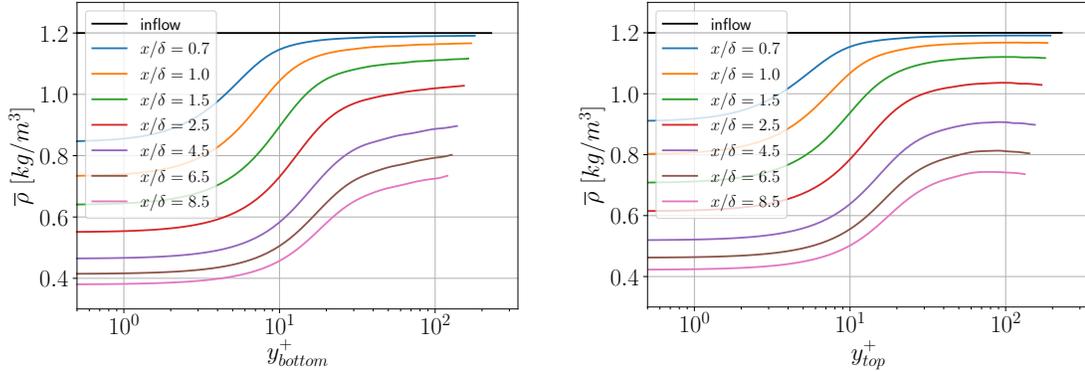
FIGURE 4.5: Streamwise development of fluid density on the bottom wall (left) and the top wall (right) for the case with $\phi_0 = 1$. Note: streamwise stations are not equally spaced in order to emphasize the sharp transition for $x/\delta < 2$.

The Reynolds-averaged mass loading plotted in figure 4.7 confirms the observation of minimal streamwise development seen in figure 4.4. This is particularly true toward the centerline of the channel, but the profile of $\overline{\phi}$ does shift somewhat very near the wall, primarily in the viscous sublayer ($y^+ < 5$). From this we see that the mass loading at the wall increases slightly at first, which can be attributed to the particles' inertia. The fluid expands first, but the particle clusters respond more slowly, i.e. the decrease in the numerator of $\phi = m_p n/\rho$ lags behind the denominator, and so ϕ increases.

Figure 4.8 shows the streamwise development of the fluid velocity, scaled with the local viscous scales in the heated domain. While the fluid is accelerating, the profiles of u^+ instead decrease in magnitude with streamwise development. This indicates that the friction velocity is increasing even more rapidly than the centerline velocity. However, for $x/\delta > 2.5$, u^+ changes very little with streamwise development. This indicates that a new balance has been reached between the increasing wall stress, expansion of the near-wall fluid, and the pressure gradient. We refer to this region as "quasi-equilibrium", or "fully developed". As expected based on other quantities, these changes are slightly more pronounced at the bottom wall. Regardless of the scaling or the side of the channel, the fluid velocity profile has become much flatter with streamwise development due to the greater expansion and acceleration of the fluid near the wall.

4.3.3 Turbulence modulation

Next, we turn to turbulent quantities. The streamwise velocity fluctuations are shown in figure 4.9 for bottom and top walls. Like the profiles of u^+, the streamwise fluctuation intensity decreases with streamwise development due to the increase in friction velocity. However, a quasi-steady state is not reached, indicating that the fluctuations themselves are also decreasing in intensity. This can
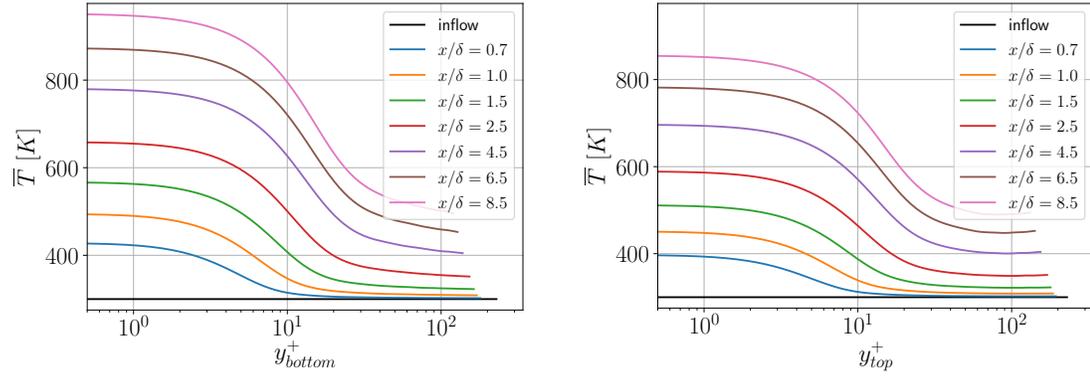
FIGURE 4.6: Streamwise development of fluid temperature on the bottom wall (left) and the top wall (right) for the case with $\phi_0 = 1$. Note: streamwise stations are not equally spaced in order to emphasize the sharp transition for $x/\delta < 2$.
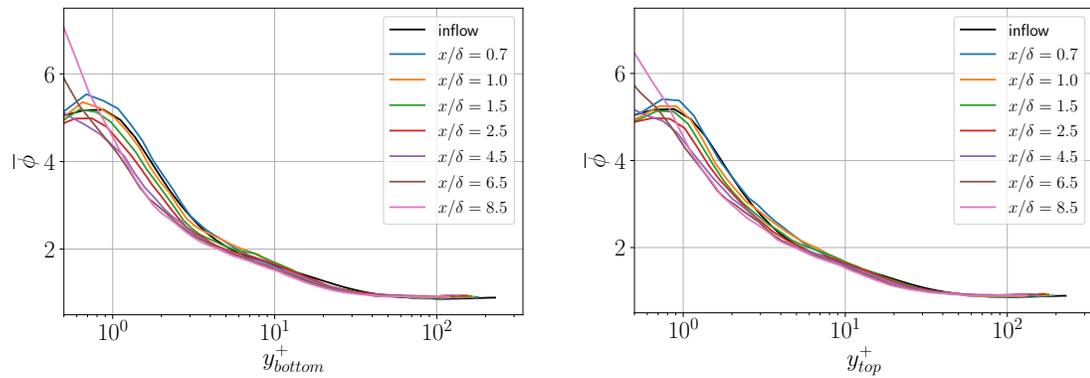
FIGURE 4.7: Streamwise development of mass loading on the bottom wall (left) and the top wall (right) for the case with $\phi_0 = 1$. Note: streamwise stations are not equally spaced in order to emphasize the sharp transition for $x/\delta < 2$.
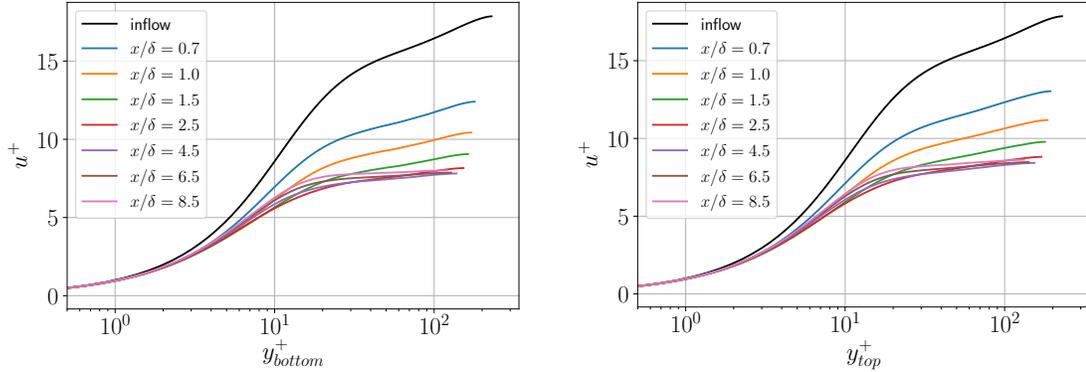
FIGURE 4.8: Streamwise development of streamwise fluid velocity on the bottom wall (left) and the top wall (right) for the case with $\phi_0 = 1$. Note: the fluid velocity is scaled with the viscous scales based on the wall stress and fluid properties at the wall in the heated domain, which all change with streamwise station.

be attributed to both the viscosity increase due to heating, which effectively lowers the Reynolds number, as well as the acceleration, which on its own can cause re-laminarization [141]. The same attenuation is observed in the wall-normal velocity fluctuations in figure 4.10. The attenuation of all velocity fluctuations is slightly weaker on the top wall, but the optical depth of the present simulations is too small for this difference to be substantial.

Next, we turn our attention to turbulent heat fluxes to assess the importance of turbulence to heat transfer in this flow. Because of the low-Mach approximation, there is a simple relationship between the turbulent heat flux and the turbulent mass flux:

$$\overline{\rho}\widetilde{u_i''T''} = -\frac{P_0}{R_g}\frac{\overline{\rho'u_i'}}{\overline{\rho}}, \tag{4.18}$$

where the $\widetilde{}$ denotes Favre averaging: $f = \widetilde{f} + f''$; $\widetilde{f} = \overline{\rho f}/\overline{\rho}$. This relationship implies that the turbulent heat fluxes are only important if the turbulent mass flux is significant relative to the mean mass flux. The turbulent heat fluxes in the streamwise and is shown in figure 4.11. The flux is negative, as expected because faster moving fluid (i.e. from the channel center) also tends to be colder, and vice versa. Unlikely the velocity fluctuations, the trend in turbulent heat flux is non-monotonic. There are no temperature fluctuations in the inflow, and so at first the turbulent heat flux increases in magnitude, reaching a peak value at approximately the same location that the velocity profile reaches its quasi-equilibrium state $x/\delta \approx 2$. After that, the heat flux begins to decay, following a similar pattern as the fluid velocity fluctuations. This is expected, because the same eddies are responsible for mixing momentum and thermal energy. Finally, we note that the streamwise turbulent heat flux is small compared to the mean heat flux, which in the viscous
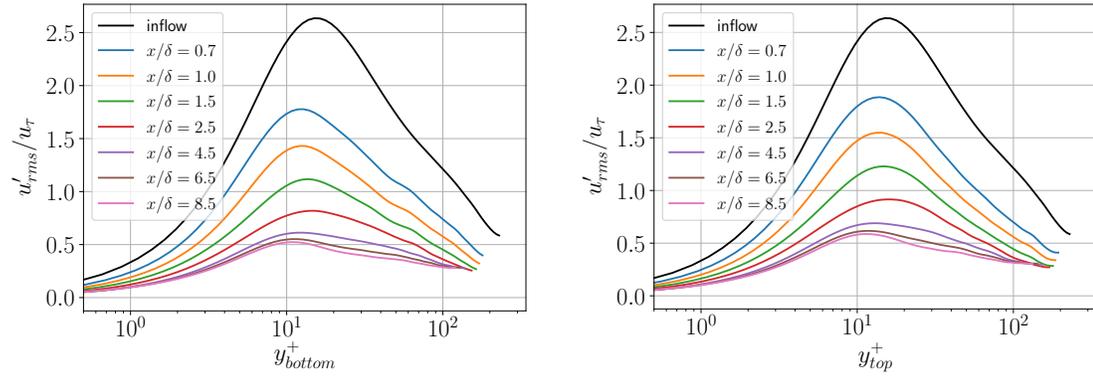
FIGURE 4.9: Streamwise development of streamwise velocity fluctuations on the bottom wall (left) and the top wall (right) for the case with $\phi_0 = 1$.
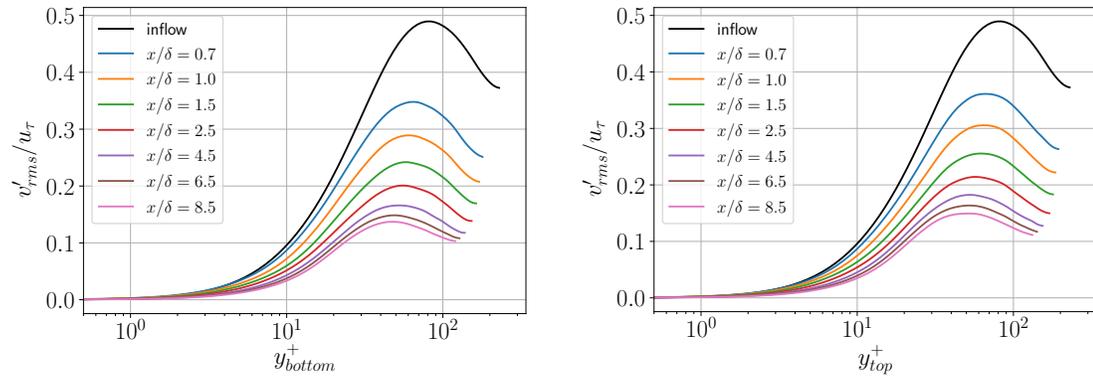
FIGURE 4.10: Streamwise development of wall-normal velocity fluctuations on the bottom wall (left) and the top wall (right) for the case with $\phi_0 = 1$.
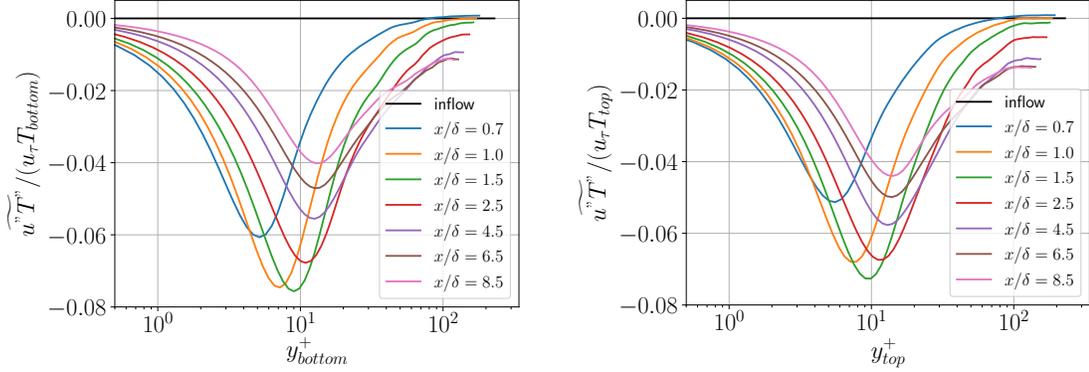
FIGURE 4.11: Streamwise development of streamwise turbulent heat flux on the bottom wall (left) and the top wall (right) for the case with $\phi_0 = 1$.
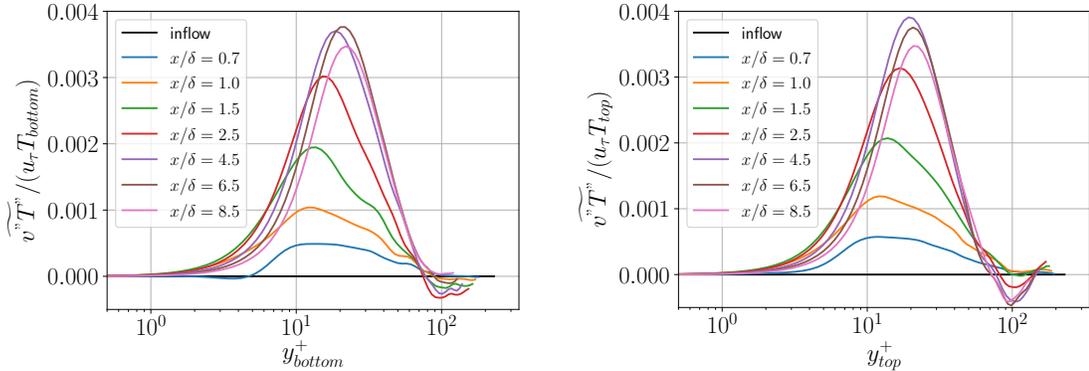
FIGURE 4.12: Streamwise development of streamwise turbulent heat flux on the bottom wall (left) and the top wall (right) for the case with $\phi_0 = 1$.

sublayer can be approximated as $\overline{u}\overline{T}/(u_\tau T_{bottom}) \approx 1$. Thus, the turbulent heat flux is only a few percent of the mean heat flux, which is dominated by the strong radiation heating.

Finally, the wall-normal turbulent heat flux is shown in figure 4.12. This heat flux is an order of magnitude smaller than the streamwise turbulent heat flux. Interestingly, this heat flux develops more slowly than the streamwise mass flux and peaks slightly later, but it eventually decreases as well, as expected from the attenuation of velocity fluctuations.

All together, the changes to instantaneous snapshots of fluid density and velocity, as well as turbulent statistics paint a picture of how strong radiation transfer can attenuate turbulence. The length of the heated part of the channel in the present simulations is short enough that the outflow is still turbulent, but for longer ducts, stronger radiation heating, or lower Reynolds numbers of the inflow, re-laminarization is increasingly likely.

4.3.4 Scaling and transforms

Figure 4.8 showed that the traditional viscous scales can effectively collapse the fluid velocity profiles to a new quasi-equilibrium state in the heated and accelerating channel flow, for the case of $\phi_0 = 1$. However, this does not collapse with the unheated inflow data. Since this flow involves not just variable mixture density (i.e. $\rho_m = \rho + m_p n$) as described in Chapter 2, but also variable fluid density and fluid viscosity, it is reasonable to expect that variable-property transforms, such as those developed by [198] and Patel et al. [158] might do a better job collapsing the velocity data to the universal incompressible profile. Accordingly, the fluid velocity profiles are plotted with three different transformations in figure 4.13: the traditional viscous scales, the mixture Trettel-Larsson transform from Chapter 2, and the mixture Trettel-Larsson transform with empirical offsets for the unheated flow, also from Chapter 2. The traditional viscous scales do a reasonably good job collapsing the quasi-steady profiles within each simulation case. However, these do not collapse to their respective inflows, nor to each other, as the streamwise acceleration (and as a consequence the wall stress) depends on the mass loading. Additionally, we note that the collapse is better for the lower mass loading cases, especially $\phi = 0.1$. If the mixture Trettel-Larsson transform is applied, the inflow data shows the same offsets observed in Chapter 2, but the collapse in the viscous sublayer across all cases has improved substantially, with all profiles being reasonably well-collapsed for $Y_m^* < 5$. This can likely be attributed to the fact that the Trettel-Larsson transform accounts for wall-normal variation in density and viscosity. The heated cases appear to collapse more closely to each other and to the universal profile than before in the log layer, but this "success" is dubious, because the positive vertical shift due to mass loading and Stokes number may be canceling out the negative vertical shift due to acceleration. Finally, we apply the empirical offsets to the mixture Trettel-Larsson transform obtained in Chapter 2 to the heated data, although we re-scale them based on the local viscous scales from the heated domain. This does not result in collapse either, although most of the remaining misalignment of profiles is presumably due to differences in acceleration caused by mass loading. We also note that for the higher mass loadings, particular $\phi_0 \geq 1$, the collapse within each simulation case has worsened, compared to the transform without offsets. This may be due to additional thickening of the viscous sublayer in these cases, which occurs due to strong viscosity increase. The offsets from Chapter 2 do not account for this mechanism of viscous sublayer thickening, and so it expected that these offsets are under-estimates.

4.4 One-dimensional model for engineering design

While the simulations described in sections 4.2 and 4.3 can be used to model the full physics of the system, they are not practical from an engineering point of view. In this section, the performance of a simplified 1D model from Banko [17] and Frankel [65] is assessed against the simulation data.
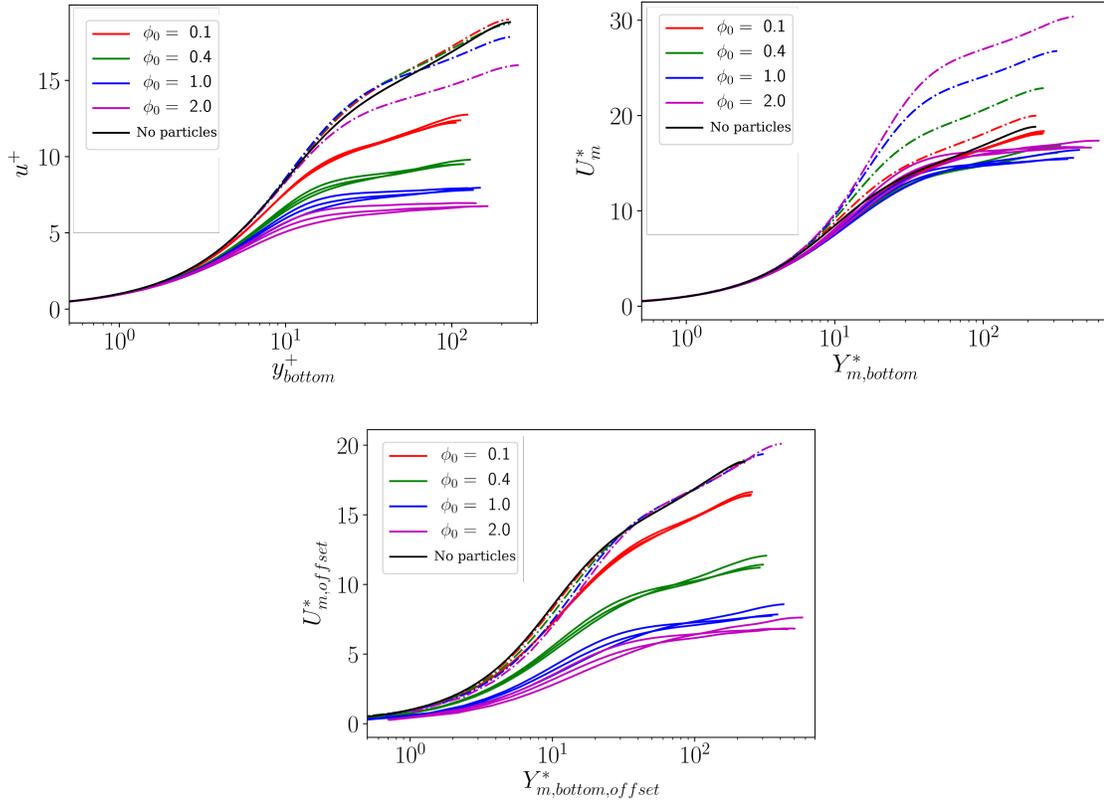
FIGURE 4.13: Streamwise velocity across mass loading variation using the traditional viscous scales for normalization (top left), the "mixture Trettel-Larsson" transform (top right), and the "mixture Trettel-Larsson" transform with empirical offset values from Chapter 2 (bottom). Unheated inflow is designated with a dash-dot line. Solid lines are from multiple stations in the "quasi-equilibrium" region at $x/\delta = [3.1, 5.2, 7.3]$.

4.4.1   Derivation

A 1D model for the heated channel flow can be obtained by assuming that there is no temporal, spanwise, or wall-normal variation in any flow quantity, i.e. $u(x, y, z, t) = u(x)$ for all flow variables. This amounts to saying that the flow is perfectly mixed at each streamwise station. Under that assumption, the mass equations for fluid and particle phase become:

$$\frac{d}{dx}(\rho u) = 0 \tag{4.19}$$

$$\frac{d}{dx}(nu_p) = 0 \tag{4.20}$$

These equations can be solved in terms of their initial values.

$$\rho u = \rho_0 u_0 \tag{4.21}$$

$$nu_p = n_0 u_{p,0} \tag{4.22}$$

Under the simplifying assumptions, the ODE for particle momentum (equation 2.5) becomes

$$\frac{du_p}{dx} = \frac{1}{\tau_p}\frac{u - u_p}{u_p}. \tag{4.23}$$

To solve this evolution equation for u_p, we need τ_p, and u. τ_p depends on the fluid temperature through the fluid viscosity. Fluid temperature can be related to fluid density through the equation of state, and density can be related to u through conservation of mass. As a result, we need an equation to obtain the fluid temperature. For this, we use the Eulerian thermal energy transport equation. With our simplifying assumptions, this becomes:

$$C_p \rho_0 u_0 \frac{dT}{dx} = \frac{d}{dx}(k\frac{dT}{dx}) - \pi D_p^2 h_p n(T - T_p) \tag{4.24}$$

Ignoring heat diffusion in the streamwise direction, which we expect to be small, this becomes

$$\frac{dT}{dx} = \frac{1}{C_p \rho_0 u_0} \pi D_p^2 h_p n(T_p - T) \tag{4.25}$$

We note that $h_p = h_p(T)$. To close this equation, we need the particle temperature, T_p, which is obtained by simplifying the ODE for particle temperature (equation 4.3)

$$\frac{dT_p}{dx} = \frac{1}{m_p C_{vp} u_p}\left(\pi D_p^2 h_p(T - T_p) + Q_{abs}\frac{\pi}{4}D_p^2 I\right). \tag{4.26}$$

To close this equation, we need the radiation intensity $I(x)$. Naively, we could use the simulation boundary condition: $I(x) = I_0$, but it is better to account for the radiation attenuation by particles,

which lowers the average radiation intensity across the channel. This can still be obtained in a 1D model by calculating the average radiation intensity

$$I_{thick}(x) = \frac{1}{2\delta} \int_0^{2\delta} I(x, y) dy. \tag{4.27}$$

To compute $I(x, y)$, we turn to the radiative transfer equation:

$$\frac{dI}{dy} = -\sigma_a I \tag{4.28}$$

which has the solution

$$I(x, y) = I_0 \exp\left(-\int_0^y \sigma_a(x, y') dy'\right) \tag{4.29}$$

Using this result and plugging in the expression for σ_a:

$$I_{thick}(x) = \frac{1}{2\delta} \int_0^{2\delta} \left(I_0 \exp\left(-\frac{\pi}{4} D_p^2 Q_{abs} \int_0^y n(x, y') dy'\right)\right) dy \tag{4.30}$$

To continue, we assume that the particles are uniformly distributed in the channel $n(x, y) = n(x)$ and simplify to obtain

$$I_{thick}(x) = I_0 \frac{1}{\frac{\pi}{2} D_p^2 Q_{abs} n\delta} \left(1 - \exp\left(-\frac{\pi}{2} D_p^2 Q_{abs} n\delta\right)\right) \tag{4.31}$$

$$\tag{4.32}$$

With this expression for average radiation intensity, the 1D model can be made to account for radiation attenuation without increasing computational cost. To summarize, the model to obtain u, u_p, T, T_p, ρ, and n is

$$u = \frac{\rho_0 u_0}{\rho}, \tag{4.33}$$

$$\frac{du_p}{dx} = \frac{1}{\tau_p} \frac{u - u_p}{u_p}, \tag{4.34}$$

$$\frac{dT}{dx} = \frac{1}{C_p \rho_0 u_0} \pi D_p^2 h_p n (T_p - T), \tag{4.35}$$

$$\frac{dT_p}{dx} = \frac{1}{m_p C_{vp} u_p} \left(\pi D_p^2 h_p (T - T_p) + Q_{abs} \frac{\pi}{4} D_p^2 I_{thick}\right), \tag{4.36}$$

$$\rho = \frac{\rho_0 T_0}{T}, \tag{4.37}$$

$$n = \frac{n_0 u_{p,0}}{u_p}. \tag{4.38}$$

This can be solved by integrating the above ODEs, given a known inflow condition based on fluid and particle mass flux and temperatures.

4.4.2   Model performance

The coupled ODEs in the 1D model are solved numerically using the odeint function from the scipy.integrate package, with default solution parameters. To compare the model predictions to the high-fidelity simulation data, the simulation data must be averaged in the spanwise and wall-normal directions. We use a bulk average based on fluid or particle mass, as appropriate. For example, the averaging formulas for fluid and particle velocities are defined as

$$U_f(x) = \frac{\int_0^{2\delta} \overline{\rho u}(x,y) dy}{\int_0^{2\delta} \overline{\rho}(x,y) dy} \tag{4.39}$$

$$U_p(x) = \frac{\int_0^{2\delta} \overline{n u_p}(x,y) dy}{\int_0^{2\delta} \overline{n}(x,y) dy} \tag{4.40}$$

The predictions of the 1D optically thick model compared to the high-fidelity simulation data are shown in figure 4.14 for fluid and particle temperatures, figure 4.15 for fluid and particle velocities, and figure 4.16 for fluid density and mass loading. Overall, the fluid quantities are predicted quite well while the particle quantities show qualitatively correct trends, and all predictions improve as mass loading increases. Figure 4.14 shows that the 1D model tends to under-predict both fluid and particle temperatures, although it under-predicts particle temperatures more, and fluid temperatures are predicted accurately from $\phi_0 = 2$. Since both are under-predicted, this suggests that the heat source is under-estimated. A possible explanation is that because particles accumulate near the wall, they have a longer residence time on average than the fluid. Since the model assumes uniform flow, this longer residence time of the hottest particles is not accounted for. This would also explain why the bulk particle temperature is under-predicted by the model. This hypothesis also fits the observed improvement as mass loading increases, because the particle distribution becomes more uniform across the channel as mass loading increases. One feature of the simulation data that the model captures qualitatively is the drop in particle temperatures after leaving the irradiated section. The particles relax back toward the fluid temperature due convective exchange, without the radiation source term acting to maintain the temperature difference between the phases.

Since the fluid temperature is under-predicted by the model, it follows that density will be over-predicted, and so from conservation of mass, it is expected that the fluid velocity will be under-predicted, as shown in figure 4.15. Again, the model predictions improve as mass loading is increased. Unlike the particle temperature, the particle velocities are over-predicted. This may be for the same reason however, because particles are more likely to be located near the wall, where the flow is both hotter and slower.
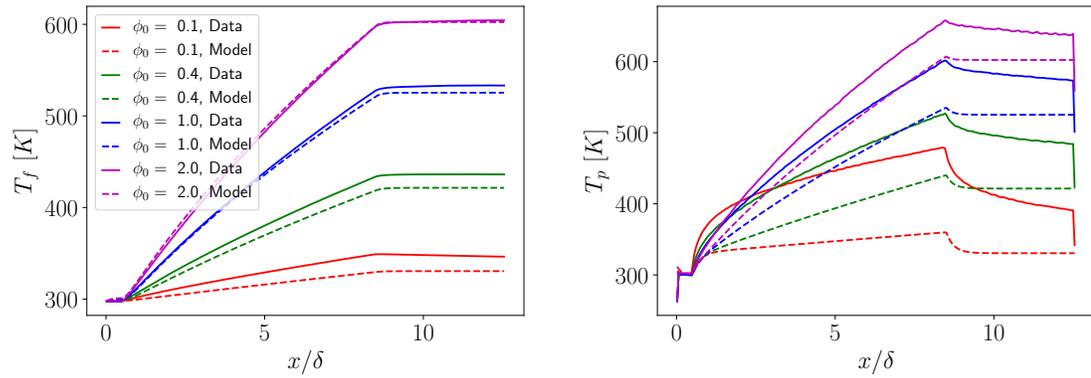
FIGURE 4.14: Average fluid temperature (left) and average particle temperature (right) obtained via simulation data and quasi-1D model.
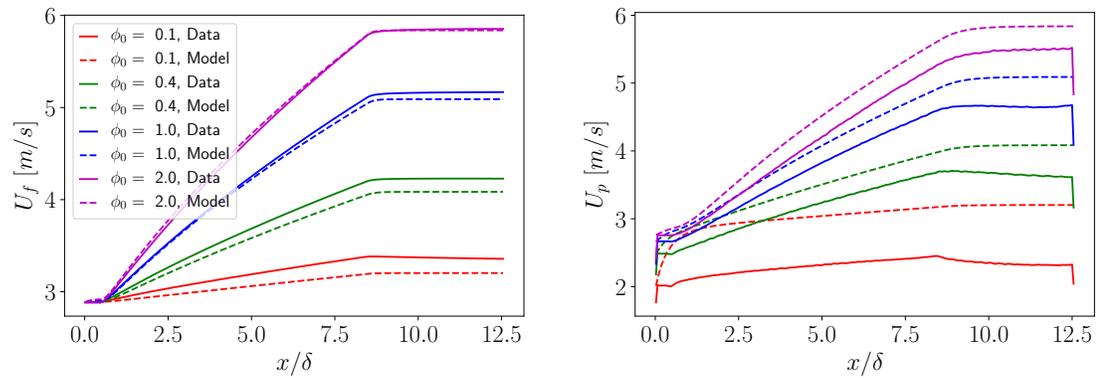
FIGURE 4.15: Average fluid velocity (left) and average particle velocity (right) obtained via simulation data and quasi-1D model.
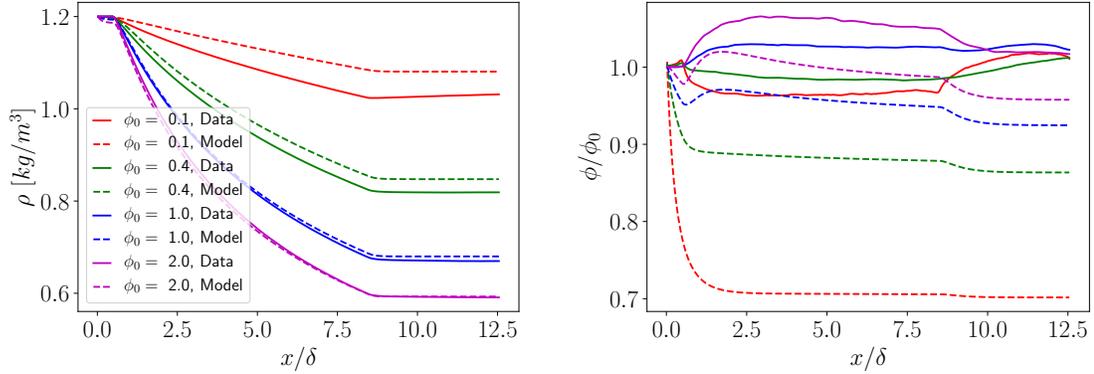
FIGURE 4.16: Average fluid density (left) and particle mass loading (right) obtained via simulation data and quasi-1D model.

Finally, the fluid density and mass loading are shown in figure 4.16. As expected, ρ is over-predicted and becomes more accurate as mass loading increases. The mass loading normalized by the averaged unheated value (ϕ_0) is shown to facilitate comparison among the cases. This shows that in the high-fidelity simulations, the mass loading is more constant with streamwise development than the model predicts, particularly for small mass loadings.

4.5 Conclusion

In this Chapter, a subset of the unheated flow simulations from Chapter 2 have been subjected to intense radiation, of a similar magnitude as would be expected in a concentrating solar power plant. The overall flow phenomena induced by radiation heating and heat transfer between the fluid and particles are expansion and acceleration of both fluid and particle phases, while the mass loading remains nearly constant with streamwise development. Because of turbophoresis, the concentration of particles is elevated near the walls, which causes much greater heat absorption and transfer to the fluid near the walls than in the centerline. Because the particles also accumulate in the low-speed streaks, this preferentially deposits heat in the streaks, and so the low-speed streaks become low-density streaks and accelerate more. Due to radiation absorption, the bottom and top walls develop differently, with less heat absorption, expansion, and acceleration on the top wall.

Immediately after the radiation boundary condition is applied, the near-wall flow responds rapidly, but the core does not. This adjustment region has a short streamwise extent, and afterwards, there is a quasi-equilibrium balance between flow acceleration and increasing wall shear stress. This results in near-collapse of velocity profiles within each mass loading case, but not with each other, because the accelerations differ. The velocity profiles become much flatter as mass loading increases. In addition to increasing flow uniformity, the combination of expansion and viscosity

increase results in damping of all turbulent fluctuations with streamwise development, with greater damping at higher mass loading. These factors together suggest that the turbulence can be killed by strong radiation heating, depending on the initial Reynolds number and degree of heat absorption.

Applying the mixture Trettel-Larsson transform to the heated data results in improved collapse in the viscous sublayer due to accounting for variable properties, but because the transform does not account for flow acceleration, collapse of log layer slopes is not obtained. When the offsets from the unheated flow are also applied to the mixture Trettel-Larsson transform, the data at large mass loading collapses worse than before, suggesting that offsets need to increase in the heated domain to account for the viscous sublayer thickening due to viscosity increase.

Finally, we have demonstrated that the 1D, optically thick model from Frankel [65] and Banko [17] is effective in predicting bulk fluid quantities which would be of interest for engineering design in the regime studied here. It is less accurate for particle quantities, because the simplifying assumptions neglect the significant inhomogeneity in the wall-normal direction due to turbophoresis. This also explains why the predictions become better as mass loading increases, since the particle distribution becomes more homogeneous in this regime.

Chapter 5

Introduction: Methods for robust simulation of shock-driven problems in elastic-plastic solids with realistic material phenomena

As described in chapter 1, multiphase flows can be either dispersed flows, such as the particle-laden flow studied in the preceding chapters, or separated flows [32]. The following chapters in this dissertation pertain to numerical methods to simulate separated multiphase flows. The methods are developed with certain problems in mind, in which there are shocks, material interfaces, and large deformations in metals. An application containing all of these features is impact welding, which involves joining together two pieces of metal impulsively, typically driven together at an oblique angle by an explosion. At the interface between the two metals, fluid-like Kelvin-Helmholtz instabilities develop when the materials yield. An example of the resulting deformation is shown in Figure 5.1.

Another application which involves shocks interacting with solid and fluid phases is inertial confinement fusion (ICF), which is visualized schematically in Figure 5.2. In ICF, laser energy is focused onto a small target, which contains deuterium gas encased in deuterium ice, further encased in a solid shell, called an ablator. The ablator implodes inward due to the intense energy deposition, compressing the materials inside. During this compression, shock waves travel through solid and fluid interfaces, and Rayleigh-Taylor and Richtmyer Meshkov instabilities are formed [46].

Simulating problems such as inertial confinement fusion and impact welding requires a numerical method capable of handling shocks and material interfaces, as well as the different stress-strain behaviors of fluids and solids. The numerical method presented here address these requirements
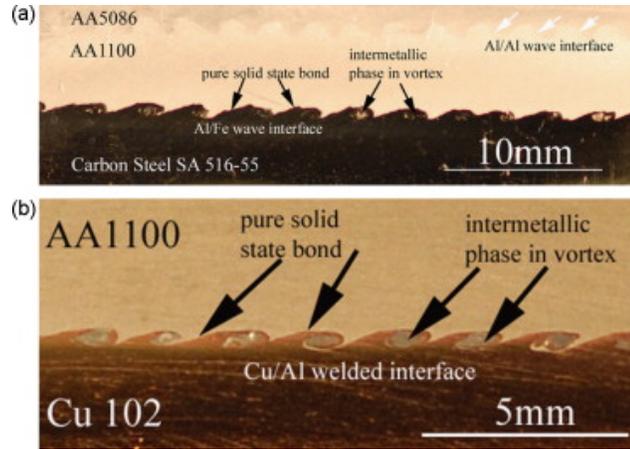
FIGURE 5.1: Kelvin-Helmoltz instabilities formed during impact welding of metals. *Image credit: Zhang et al. [217]*
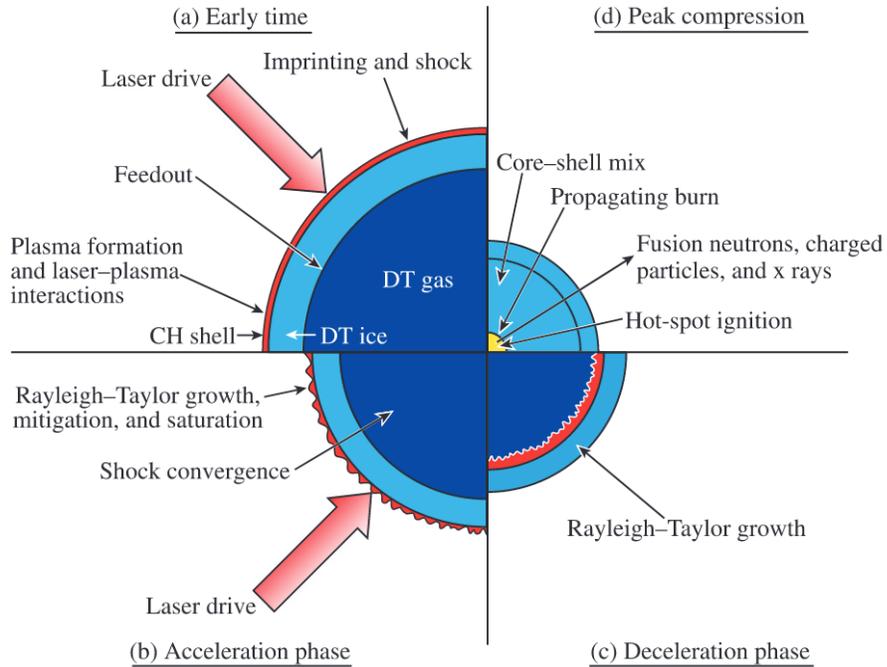
FIGURE 5.2: Schematic of direct drive inertial confinement fusion. Sub-figures (a), (b), (c), and (d) show phases of an ICF experiment in chronological order. *Image credit: Craxton et al. [46]*

by building on prior work by Subramaniam, Ghaisas, and Lele, in particular [73] and [188]. The following chapters extend their method to simulate problems involving strain hardening (Chapter 6) and interface-sharpening (Chapter 7).

In Chapter 6, a high-order method for Eulerian simulation of material undergoing large elastic-plastic deformation is developed. Thermodynamically consistent hyperelastic constitutive relations are assumed, facilitating the treatment of solids, liquids, and gases in a unified manner. The method enables the simulation of multi-material interactions using a diffuse interface approach. Numerical capturing of material interfaces, shock waves, contact surfaces, and elastic-plastic strain disconti-nuities using high-order compact-difference schemes is assisted by Localized Artificial Diffusivity (LAD). Previously used terms for artificial properties are verified to appropriately regularize normal shocks, whether they involve elastic or plastic deformation. Additional LAD terms are added to elastic and plastic kinematic equations to regularize shear shocks and other strain discontinuities, and to improve solution stability. Other important features of the method which improve its ro-bustness include the numerical treatment of compatibility terms in the kinematic equations, and the treatment of rotation. Particular emphasis is focused toward new advancements of the meth-ods for plastic-deformation integration and associated strain hardening of the material, including rate-dependent plasticity. The method is demonstrated on a variety of test problems, including 1d impacts, a variant of the Shu-Osher problem, a Taylor impact problem and a Richtmyer-Meshkov instability between two elastic-plastic solids with strain hardening.

In the Richtmyer Meshkov results, it becomes clear that the interface treatment has some down-sides, namely that the localized artificial diffusivity approach can only make interfaces more diffuse, although that diffusion is localized and limited. This deficiency shows up in regions of under-resolved interface roll-up, and results in unphysical mixing between solid phases. One way of addressing this problem is with interface sharpening techniques, three of which are considered in Chapter 7. The first method is the *localized-artificial-diffusivity approach* of Cook [39], Subramaniam et al. [188], and Adler & Lele [5], in which artificial diffusion terms are added to the individual phase mass frac-tion transport equations and are coupled with the other conservation equations. The second method is the *gradient-form approach* that is based on the quasi-conservative method of Shukla et al. [180], in which the diffusion and sharpening terms (together called regularization terms) are added to the individual phase volume fraction transport equations and are coupled with the other conservation equations [196]. The third approach is the *divergence-form approach* that is based on the fully con-servative method of Jain et al. [97], in which the regularization terms are added to the individual phase volume fraction transport equations and are coupled with the other conservation equations. The primary objective of this chapter is to compare these three methods in terms of their ability to: maintain constant interface thickness throughout the simulation; conserve mass, momentum, and energy; and maintain accurate interface shape for long-time integration. The second objective is to consistently extend these methods to model interfaces between solid materials with strength. To

assess and compare the methods, they are used to simulate a wide variety of problems, including (1) advection of an air bubble in water, (2) shock interaction with a helium bubble in air, (3) shock interaction and the collapse of an air bubble in water, and (4) Richtmyer–Meshkov instability of a copper–aluminum interface. The current work focuses on comparing these methods in the limit of relatively coarse grid resolution, which illustrates the true performance of these methods. This is because it is rarely practical to use hundreds of grid points to resolve a single bubble or drop in large-scale simulations of engineering interest.

Together, Chapters 6 and 7 advance the state of Eulerian finite difference methods for shocks in elastic-plastic solids by extending existing high-order, localized artificial diffusivity techniques to problems involving strain hardening, and by extending explicit interface sharpening methods to problems involving solids.

Chapter 6

A high-order, localized artificial diffusivity method for Eulerian simulation of multi-material elastic-plastic deformation with strain hardening

6.1 Introduction

Shock-wave propagation through multi-material media is relevant to physical problems such as detonations, impact welding, and ejecta formation. Recent interest in the improvement of accuracy for the simulation of this phenomenon in elastic-plastic materials has been motivated by ICF efforts, which require extreme shock compression of elastic-plastic material, namely the ablator [118]. In ICF experiments, the development of hydrodynamic instabilities, including Richtmyer-Meshkov (RMI) and Rayleigh-Taylor, has been observed [220]. Material strength is of interest as a means of suppressing these instabilities and increasing the compression efficiency, and simulations are relied upon to target experimental campaigns. Other relevant physical problems include impact welding [142], in which instability growth between metal interfaces may be important to weld strength. Additionally, RMI can be used to determine material properties under extreme deformation conditions [51]. Experimental and computational techniques must advance together to improve engineering design capabilities. In solid materials, shock waves can create large deformations at high strain

rates, resulting in strain hardening, thermal softening, and other inelastic phenomena. This requires a numerical method which can accurately resolve instability growth, material interfaces, and elastic-plastic deformation with complex material behavior in a unified framework. In multi-phase problems, this may also require resolving shock-turbulence interaction [3, 4].

A variety of approaches can be used to simulate material deformation, including Lagrangian [26], arbitrary Lagrangian-Eulerian (ALE) [52], and fully Eulerian methods. In the ICF application, the material interface seeds the hydrodynamic instability and undergoes severe distortions in the transitional or turbulent regimes. Similar difficulties occur in the impact welding problem, where fluid-like instabilities develop at a metal-metal interface. In cases of severe interface distortions, the Lagrangian approach can result in mesh entanglement. The ALE approach can mitigate this issue by rezoning and remapping the mesh, but this can introduce additional instability, especially when the solution contains shocks [83]. To avoid these types of errors, and to leverage previous development in localized artificial diffusivity methods, we opt for a fully Eulerian approach.

Eulerian methods can represent interfaces numerically as sharp or diffuse [136, 175]. In the ICF problem, the strongly conservative diffuse-interface approach presents an advantage over the sharp-interface approach. Because the time scales of the relevant deformation are short relative to the time scales of numerical diffusion of the diffuse interface, no interface sharpening procedure is needed, although existing interface sharpening methods can be adapted to solids [98].

Seminal theoretical developments of fully Eulerian diffuse-interface approaches for elastic-plastic deformation are described by Godunov and Romenskii [75], Plohr and Sharp [161], and Trangenstein and Colella [197]. A few recent applications of these approaches are presented by Hank et al. [81], Ndanou et al. [144], and Ortega et al. [148], who have demonstrated the success of the method for a variety of high-velocity solid–solid and solid–fluid impact scenarios. Several areas of improvement are still desired, including improved physical aspects such as the inclusion of strain and strain rate in the parameterized plasticity model, more robust treatment of thermal softening and phase changes, and improved equations of state.

Another limiting aspect of many efforts to date has been the use of relatively low-order explicit spatial schemes in conjunction with a potentially costly, complex, and dissipative multi-material Riemann solver. An alternative approach has been recently examined, employing a combination of high-order compact difference schemes and localized artificial diffusivity (LAD) [73]. Such methods have proven successful for a variety of shock-turbulence interaction problems and other complex multi-scale flows [101, 199]. This approach provides the necessary low-dissipation/high-resolution numerical scheme for accurate reproduction of fine-scale turbulence and, by means of LAD, functions as a high-dissipation/(inherently low-resolution) scheme for shock-wave and material-interface capturing. Recent extension of this method for use in diffuse-interface capturing of multi-material interfaces has demonstrated superior resolution properties relative to non-compact reconstruction schemes [188].

The present work concerns the extension of the numerical method developed by Subramaniam, Ghaisas, and Lele [73, 188] to problems involving strain-dependent plasticity. The LAD approach is extended to appropriately regularize the variety of discontinuities which can arise in elastic-plastic solids, and other improvements are made to the method to improve its robustness in challenging problems involving significant plastic shear. For many materials, plastic deformation alters the microstructural properties of the material, increasing its strength. For metals, this strain hardening can increase the ultimate strength of the material in excess of several times the pre-hardened yield strength [19]. In the context of ICF-relevant deformations, the effect of strain hardening can be significant on the development of hydrodynamic instabilities [149], suppressing the Richtmyer-Meshkov growth rate, increasingly so for smaller Mach numbers.

However, in the context of fully Eulerian diffuse-interface methods, the examination of strain-dependent plasticity has thus far been fairly limited. A few notable examples include the work of Plohr and Sharp [162], who present a theoretical framework for transport of a plastic strain tensor (6 additional equations per material) and work-hardening variable (+1 equation). Later, Miller and Colella [135] examined the use of the full plastic deformation gradient tensor (+9 equations) for strain-hardening parametrization. More recently, Barton and Romenski [19] proposed the use of a reduced formulation based on plastic entropy transport (+1 equation). These formulations span space spans a wide range of computational complexity and modeling assumptions. The optimal strategy remains unclear and certainly depends on the deformation scenario and calculation objectives regarding accuracy and computational efficiency. In this work, we present a formulation similar to Plohr and Sharp [162], which requires 6 additional equations per material for plastic deformation.

The reference map technique has also been used to compute the motion of elastic solids in multi-phase problems, with a reduced number of equations. This approach involves solving an evolution equation for a vector-valued mapping between the current and previous position of material parcels (3 equations) [104, 200]. Gradient information of this field can be used to compute elastic energies and stresses, but how to extend the method to handle elastic-plastic deformation is unclear. Early versions of the reference map formulation included an artificial viscous stress for the solid phase, which is used for quasi-steady problems in which the dynamics of elastic waves are not important. Artificial viscous stress is applied to all phases here, but in addition, artificial viscosity is also applied to the kinematic equations themselves, to regularize them when resolving shocks.

The remainder of this article is outlined as follows. The numerical framework for the Eulerian, multi-material, diffuse-interface approach, with recent advancement in strain-hardening parametrization, as well as other changes to improve robustness is described in Section 6.2. A variety of validation test problems are presented in section 6.3. Validation against other simulations and experimental results using a variety of strain hardening models is presented for a Taylor impact impact problem in Section 6.4. Finally, a demonstration of the proposed method on the Richtmyer-Meshkov instability of a 2-D multi-material interface is given in Section 6.5, including the coupled effects of strain

hardening. Concluding remarks are made in Section 6.6.

6.2 Theoretical and numerical model

6.2.1 Multi-material Eulerian conservation equations

The governing equations for the evolution of a multi-material continuum in conservative Eulerian form are reproduced below, including the conservation of species mass, total momentum, and total energy.

$$\frac{\partial \rho Y_m}{\partial t} + \nabla \cdot (\rho Y_m \boldsymbol{u}) = -\nabla \cdot (\boldsymbol{J}_m^*) \tag{6.1}$$

$$\frac{\partial \rho \boldsymbol{u}}{\partial t} + \nabla \cdot (\rho \boldsymbol{u} \otimes \boldsymbol{u} - \underline{\boldsymbol{\sigma}}) = \nabla \cdot (\underline{\boldsymbol{\tau}}^*) \tag{6.2}$$

$$\frac{\partial}{\partial t}\left[\rho\left(e + \frac{1}{2}\boldsymbol{u}\cdot\boldsymbol{u}\right)\right] + \nabla\cdot\left(\left[\rho\boldsymbol{u}\left(e + \frac{1}{2}\boldsymbol{u}\cdot\boldsymbol{u}\right) - \underline{\boldsymbol{\sigma}}\cdot\boldsymbol{u}\right]\right) = \nabla\cdot(\underline{\boldsymbol{\tau}}^*\cdot\boldsymbol{u} - \boldsymbol{q}^*) \tag{6.3}$$

Concerning notation, t and $\boldsymbol{x}$ represent time and the Eulerian position vector, respectively. Y_m describes the mass fraction of each constituent material, m. The variables $\boldsymbol{u}$, ρ, e, and $\underline{\boldsymbol{\sigma}}$ describe the mixture velocity, density, internal energy, and Cauchy stress, respectively, which are related to the species-specific values by the relations: $\rho = \sum_{m=1}^{M}\phi_m\rho_m$, $e = \sum_{m=1}^{M}Y_m e_m$, and $\underline{\boldsymbol{\sigma}} = \sum_{m=1}^{M}\phi_m\underline{\boldsymbol{\sigma}}_m$, in which ϕ_m is the volume fraction of material m, and M is the total number of material constituents. The right-hand side terms describe the localized artificial diffusion (based on Cook [39]), including the artificial viscous stress, $\underline{\boldsymbol{\tau}}^* = 2\mu^*\underline{\boldsymbol{S}} + (\beta^* - 2\mu^*/3)(\nabla\cdot\boldsymbol{u})\underline{\boldsymbol{\delta}}$ and the artificial enthalpy flux, $\boldsymbol{q}^* = -\kappa^*\nabla T + \sum_{m=1}^{M}(h_m\boldsymbol{J}_m^*)$, with strain rate tensor, $\underline{\boldsymbol{S}} = (1/2)\left(\nabla\boldsymbol{u} + \nabla\boldsymbol{u}^T\right)$, and temperature T. The second term in the artificial enthalpy flux expression is the enthalpy diffusion term [40], in which $h_m = e_m + p_m/\rho_m$ is the enthalpy of species m. The artificial Fickian diffusion of species m is described by $\boldsymbol{J}_m^* = -\rho\left[D_m^*\nabla Y_m - Y_m\sum_{j=1}^{M}\left(D_j^*\nabla Y_j\right)\right]$, which assumes a binary mixture between species m and a composite mixture of the other species.

6.2.2 Kinematic equations for elastic-plastic deformation

To compute the Cauchy stress in elastic-plastic solids, as well as elastic energy, the above system of conservation equations is augmented with a set of kinematic equations. To describe the deformation of a solid in the Eulerian framework, the inverse deformation gradient tensor is employed, $\underline{\mathbf{g}} = \partial \boldsymbol{X}/\partial \boldsymbol{x}$ ($g_{ij} = \partial X_i/\partial x_j$), in which $\boldsymbol{X}$ or X_i describe the position of a continuum parcel in the material (Lagrangian) perspective, and $\boldsymbol{x}$ or x_i describe the spatial (Eulerian) perspective. An inverse deformation gradient may be tracked for each material constituent (multiple-g formulation) or a single inverse deformation gradient may be tracked for the continuum mixture as a whole (mixture-g formulation). In many scenarios the accuracy achieved through use of the mixture-g formulation is indistinguishable from the multiple-g formulation outside of the diffuse interface

region [72, 73]. Following Miller and Colella [135], a multiplicative decomposition of the total inverse deformation gradient tensor, $\underline{\mathbf{g}}$, into elastic, $\underline{\mathbf{g}}^e$, and plastic $\underline{\mathbf{g}}^p$, components is used, $\underline{\mathbf{g}} = \underline{\mathbf{g}}^p\underline{\mathbf{g}}^e$, reflecting the assumption that the plastic deformation is recovered when the elastic deformation is reversed, $\underline{\mathbf{g}}^p = \underline{\mathbf{g}}\underline{\mathbf{g}}^{e\,-1}$.

It is additionally assumed that the plastic deformation is volume preserving [162], providing compatibility conditions for the inverse deformation gradient tensor determinants, $\left|\underline{\mathbf{g}}^p\right| = 1$ and $\left|\underline{\mathbf{g}}\right| = \left|\underline{\mathbf{g}}^e\right| = \rho/\rho_0$, in which ρ_0 represents the undeformed density. Recalling that the curl of a gradient of a vector must be zero, a curl compatibility condition (Equation 6.4) is provided for the inverse deformation gradient tensor, where ϵ_{ijk} denotes the Levi-Civita permutation tensor of order three.

$$\nabla \times \underline{\mathbf{g}}^T = \epsilon_{ijk}\frac{\partial g_{lk}}{\partial x_j} = \epsilon_{ijk}\frac{\partial}{\partial x_j}\frac{\partial}{\partial x_k}X_l = 0 \implies \frac{\partial g_{ij}}{\partial x_k} - \frac{\partial g_{ik}}{\partial x_j} = 0 \tag{6.4}$$

With this compatibility condition, often the kinematic equations for the evolution of the inverse deformation gradient tensor are rearranged from the form including the advection and strain-source terms (Equation 6.5, left) to the mathematically equivalent form including a curl-free advection/strain term and a non-zero curl advection/strain term (Equation 6.5, right); the latter term will be zero when the curl compatibility condition is satisfied. In practice, this term is small but not identically zero due to numerical error and should not be neglected so as to maintain hyperbolicity of the equations [143].

$$\underbrace{\frac{\partial \underline{\mathbf{g}}}{\partial t}}_{\substack{\text{local}\\\text{derivative}}} + \underbrace{\boldsymbol{u}\cdot\nabla\underline{\mathbf{g}}}_{\text{advection}} + \underbrace{\underline{\mathbf{g}}\nabla\boldsymbol{u}}_{\text{strain source}} = \mathbf{0} \iff \underbrace{\frac{\partial \underline{\mathbf{g}}}{\partial t}}_{\substack{\text{local}\\\text{derivative}}} + \underbrace{\nabla\left(\underline{\mathbf{g}}\cdot\boldsymbol{u}\right)}_{\substack{\text{curl-free}\\\text{advection/strain}}} + \underbrace{\boldsymbol{u}\cdot\nabla\underline{\mathbf{g}} - \left(\nabla\underline{\mathbf{g}}\right)\cdot\boldsymbol{u}}_{\substack{\text{non-zero curl}\\\text{advection/strain}}} = \mathbf{0} \tag{6.5}$$

Miller and Colella [135] describe a correction term to enforce the curl compatibility condition; however, we observe poor numerical stability of this correction using our method and instead follow Favrie and Gavrilyuk [60], Barton et al. [21], Ortega et al. [148], and Ndanou et al. [144], omitting an explicit curl compatibility correction.

With these assumptions, the kinematic equations describing the evolution of the total (Equation 6.6), elastic (Equation 6.7), and plastic (Equation 6.8) inverse deformation gradient tensors in a single material are given below.

$$\underbrace{\frac{\partial \underline{\mathbf{g}}}{\partial t}}_{\substack{\text{local}\\\text{derivative}}} + \underbrace{\nabla\left(\underline{\mathbf{g}}\cdot\boldsymbol{u}\right)}_{\substack{\text{curl-free}\\\text{advection/strain}}} + \underbrace{\boldsymbol{u}\cdot\nabla\underline{\mathbf{g}} - \left(\nabla\underline{\mathbf{g}}\right)\cdot\boldsymbol{u}}_{\substack{\text{non-zero curl}\\\text{advection/strain}}} = \mathbf{0} \tag{6.6}$$

$$\underbrace{\frac{\partial \underline{\mathbf{g}}^e}{\partial t}}_{\substack{\text{local}\\\text{derivative}}} + \underbrace{\nabla\left(\underline{\mathbf{g}}^e\cdot\boldsymbol{u}\right)}_{\substack{\text{curl-free}\\\text{advection/strain}}} + \underbrace{\boldsymbol{u}\cdot\nabla\underline{\mathbf{g}}^e - \left(\nabla\underline{\mathbf{g}}^e\right)\cdot\boldsymbol{u}}_{\substack{\text{non-zero curl}\\\text{advection/strain}}} - \underbrace{\underline{\mathbf{L}}^p\underline{\mathbf{g}}^e}_{\text{elastic-plastic source}} = \mathbf{0} \tag{6.7}$$

$$\underbrace{\frac{\partial \underline{\mathbf{g}}^p}{\partial t}}_{\substack{\text{local} \\ \text{derivative}}} + \underbrace{\boldsymbol{u} \cdot \nabla \underline{\mathbf{g}}^p}_{\text{advection}} + \underbrace{\underline{\mathbf{g}}^p \underline{\mathbf{L}}^p}_{\text{elastic-plastic source}} = \mathbf{0} \qquad (6.8)$$

The $\underline{\mathbf{g}}$ equation is pure kinematics, containing only terms for advection and strain. The $\underline{\mathbf{g}}^e$ equation includes a model term to describe plastic yielding, and the $\underline{\mathbf{g}}^p$ equation is derived as a consequence of the other two. The plasticity model is discussed in greater detail in Section 6.2.3.

In purely elastic deformation, in which the yield criterion is never exceeded, the plasticity model is never activated, making equations 6.6 and 6.7 equivalent, and equation 6.8 unnecessary. In the case of perfect plasticity, in which the material yield stress is unaffected by plastic deformation, only Equation 6.7 need be solved, as $\underline{\mathbf{g}}^e$ is sufficient to compute the elastic energy and elastic stresses which couple to the other equations. Strain-hardening models are parameterized by plastic strain and/or plastic strain rate, which necessitates solving an additional equation to obtain this information. For this purpose, any two of Equations 6.6-6.8 will suffice, as the third tensor can be obtained from the other two from the multiplicative decomposition ($\underline{\mathbf{g}} = \underline{\mathbf{g}}^p \underline{\mathbf{g}}^e$). There is precedent for solving evolution equations for two constituent tensors in the multiplicative decomposition from Miller and Colella [135], who chose to solve for the total inverse deformation gradient tensor ($\underline{\mathbf{g}}$) and plastic deformation gradient ($\underline{\mathbf{g}}^p$).

For the isotropic, hyperelastic materials of interest in this work, the Cauchy stress and equation of state may be expressed in terms of the (symmetric) elastic Finger tensor $\left(\underline{\mathbf{G}}^e = \underline{\mathbf{g}}^{eT} \underline{\mathbf{g}}^e\right)$, without requiring the solution of the full elastic inverse deformation gradient tensor. For perfect plasticity, this allows for a reduction in the number of kinematic equations solved from nine (g^e formulation) to six (G^e formulation). As a consequence, the rotation associated with $\underline{\mathbf{g}}^e$ is not recoverable from $\underline{\mathbf{G}}^e$. Similarly, for isotropic strain-hardening models, calculation of the plastic strain and strain rate do not require the solution of the full plastic inverse deformation gradient tensor, but may be achieved through the solution of the symmetric plastic Finger tensor $\left(\underline{\mathbf{G}}^p = \underline{\mathbf{g}}^{pT} \underline{\mathbf{g}}^p\right)$.

However, in pursuit of extending the Finger tensor formulation to also track plastic deformation, it is evident that the multiplicative decomposition of the inverse deformation gradient tensor into plastic and elastic components cannot be reduced to a form depending only on Finger tensor constituents. For example, $\underline{\mathbf{G}} = \underline{\mathbf{g}}^{eT} \underline{\mathbf{G}}^p \underline{\mathbf{g}}^e$ decomposes the total deformation Finger tensor into a plastic deformation Finger tensor and elastic inverse deformation gradient tensors, which cannot cannot be expressed in terms of $\underline{\mathbf{G}}^e$. Nevertheless, a simplification of one constituent of this decomposition is still useful, and the kinematic equations describing the evolution of the total (Equation 6.9), elastic (Equation 6.10), and plastic (Equation 6.11) Finger tensors are presented below. These are derived as a consequences of Equations 6.6-6.8. Notably, the elastic-plastic source term in Equation 6.11 depends on $\underline{\mathbf{g}}^e$ not $\underline{\mathbf{G}}^e$, reflecting the irreducibility of the Finger tensor decomposition. The G^e formulation has been discussed by several authors [75, 148, 71] and implemented in a limited number of studies [73, 133, 195]. Implementation and evaluation of the $\underline{\mathbf{G}}^p$ equations appears unique to the

present study.

$$\underbrace{\frac{\partial \underline{\mathbf{G}}}{\partial t}}_{\substack{\text{local} \\ \text{derivative}}} + \underbrace{\boldsymbol{u} \cdot \nabla \underline{\mathbf{G}}}_{\text{advection}} + \underbrace{\underline{\mathbf{G}} \nabla \boldsymbol{u} + \left(\nabla \boldsymbol{u}^T \right) \underline{\mathbf{G}}}_{\text{strain source}} = \mathbf{0} \tag{6.9}$$

$$\underbrace{\frac{\partial \underline{\mathbf{G}}^e}{\partial t}}_{\substack{\text{local} \\ \text{derivative}}} + \underbrace{\boldsymbol{u} \cdot \nabla \underline{\mathbf{G}}^e}_{\text{advection}} + \underbrace{\underline{\mathbf{G}}^e \nabla \boldsymbol{u} + \left(\nabla \boldsymbol{u}^T \right) \underline{\mathbf{G}}^e}_{\text{strain source}} - \underbrace{\underline{\mathbf{g}}^{eT} \left(\underline{\mathbf{L}}^p + \underline{\mathbf{L}}^{pT} \right) \underline{\mathbf{g}}^e}_{\text{elastic-plastic source}} = \mathbf{0} \tag{6.10}$$

$$\underbrace{\frac{\partial \underline{\mathbf{G}}^p}{\partial t}}_{\substack{\text{local} \\ \text{derivative}}} + \underbrace{\boldsymbol{u} \cdot \nabla \underline{\mathbf{G}}^p}_{\text{advection}} + \underbrace{\underline{\mathbf{L}}^{pT} \underline{\mathbf{G}}^p + \underline{\mathbf{G}}^p \underline{\mathbf{L}}^p}_{\text{elastic-plastic source}} = \mathbf{0} \tag{6.11}$$

In the case of perfect plasticity, Equation 6.10 (G^e formulation) may be solved instead of Equation 6.7 (g^e formulation), reducing the number of kinematic equations from nine to six. Equation 6.10 appears to need all 9 components of the $\underline{\mathbf{g}}^e$, but because of the specific form of $\underline{\mathbf{L}}^p$ (described in section 6.2.3), this can be expressed in terms of just $\underline{\mathbf{G}}^e$. In the case of realistic strain hardening, Equation 6.7 must be solved in addition to either Equation 6.9 (G–g^e formulation) or Equation 6.11 (G^p–g^e formulation). Both of these strain hardening formulations require the solution to three fewer equations than their full tensor counterparts.

Hill et al. [82] discuss an evolution equation for the specific entropy associated with plastic deformation, which others have explored further [19, 21, 147]. For comparison with the above formulations, we reproduce it here:

$$\frac{\partial \rho S^p}{\partial t} + \nabla \cdot (\rho S^p \boldsymbol{u}) = \frac{1}{T} \left[L^{S^p} + \underline{\boldsymbol{\tau}}^* : \nabla \boldsymbol{u} - \nabla \cdot (\boldsymbol{q}^*) \right] \tag{6.12}$$

An alternative reduced-dimension formulation for the kinematic equations involves the solution of Eq. 6.7 along with a simpler one-variable equation for plastic entropy, as described by Eq. 6.12 (S^p–g^e formulation). Here, S^p represents the entropy generated by plastic deformation and L^{S^p} represents the plastic entropy source term. This reduced-dimension formulation is based on the work of Barton and Romenski [19], which has since been discussed by Ortega et al. [149] and Ghaisas et al. [73]. In previous work, we found that the reduced formulation involving plastic entropy was accurate for normal shock deformation, but did not provide accurate calculation of equivalent plastic strain in problems involving deformation at interfacial instabilities between metals [5].

6.2.3 Plasticity model

Plastic deformation is incorporated into the numerical framework by means of a visco-plastic Maxwell relaxation model from Favrie and Gavrilyuk [59], which has been employed recently in several

Eulerian approaches [144, 149, 73]. The plastic source terms are described by

$$\underline{\mathbf{L}}^p = \frac{1}{2\mu\tau_{rel}}\underline{\mathbf{g}}^e\underline{\boldsymbol{\sigma}}'\underline{\mathbf{g}}^{e-1}; \qquad L^{S^p} = \frac{\underline{\boldsymbol{\sigma}}' : \underline{\boldsymbol{\sigma}}'}{2\mu\tau_{rel}},$$ (6.13)

and the plastic relaxation time scale is described by

$$\frac{1}{\tau_{\text{rel}}} = \frac{1}{(\rho/\rho_0)\tau_0}\left[\frac{\text{R}\left(\underline{\boldsymbol{\sigma}}' : \underline{\boldsymbol{\sigma}}' - \frac{2}{3}\sigma_Y^2\right)}{\mu^2}\right],$$ (6.14)

in which, $\underline{\boldsymbol{\sigma}}' = \text{dev}(\underline{\boldsymbol{\sigma}})$ and τ_0 is a timescale associated with plastic yielding. The ramp function $R(x) = \max(x, 0)$ turns on plasticity effects only when the yield criterion is satisfied. This model causes the deviatoric part of the elastic stress tensor to decay when the yield criterion is exceeded.

In general the yield stress and shear modulus are functions of temperature, pressure, plastic strain, and strain rate:

$$\sigma_Y = \sigma_Y(T, p, \epsilon_p, \dot{\epsilon}_p)$$ (6.15)

$$\mu = \mu(T, p, \epsilon_p, \dot{\epsilon}_p)$$ (6.16)

These dependencies allow this plastic relaxation model to describe complex material behavior, such as strain hardening and thermal softening. In many cases, the elastic-plastic source term is stiff due to the small value of τ_{rel} relative to the convective deformation scales. To overcome this time step restriction, an implicit plastic relaxation method is used, based on the method of Favrie and Gavrilyuk [59], who describe the relaxation procedure for $\underline{\mathbf{G}}^e$. The relaxation procedure for $\underline{\mathbf{g}}^e$ is very similar and is described by Subramaniam et al. [188]. Given the inverse deformation gradient tensor $\underline{\mathbf{g}}^e$ before plastic relaxation and $\widetilde{\underline{\mathbf{g}}^e}$ after plastic relaxation, the other components of the multiplicative decomposition may be computed under the assumption that plastic yielding does not change the total deformation ($\underline{\mathbf{g}}$ and $\underline{\mathbf{G}}$), but instead changes how that deformation is partitioned into elastic and plastic components:

$$\widetilde{\underline{\mathbf{G}}} = \underline{\mathbf{G}}$$ (6.17)

$$\widetilde{\underline{\mathbf{G}}^p} = \widetilde{\underline{\mathbf{g}}^e}^{-T}\underline{\mathbf{g}}^{eT}\underline{\mathbf{G}}^p\underline{\mathbf{g}}^e\widetilde{\underline{\mathbf{g}}^e}^{-1}$$ (6.18)

Strain hardening models are commonly parameterized in terms of the equivalent plastic strain and equivalent plastic strain rate (e.g. Steinberg et al. [186] and Johnson and Cook [102]). Equivalent plastic strain (ϵ_p) is computed here from a norm of the Euler-Almansi plastic strain tensor:

$$\epsilon_p = \sqrt{\frac{2}{3}\boldsymbol{e}_p : \boldsymbol{e}_p} \;\; ; \;\; \boldsymbol{e}_p = \frac{1}{2}\left(\mathbf{1} - \underline{\mathbf{G}}^p\right)$$ (6.19)

The 2/3 factor is included to relate the general strain tensor to uniaxial strain measurements under the assumption of isochoric plastic flow. Recognizing that $\dot{\boldsymbol{e}_p} = -(1/2)\underline{\dot{\mathbf{G}}}^p$, the equivalent plastic

strain rate ($\dot{\epsilon}_p$) is approximated by using the plastic Finger tensors before and after relaxation:

$$\dot{\epsilon}_p \approx \sqrt{\frac{1}{6}\left(\frac{\widetilde{\mathbf{G}^p} - \mathbf{G}^p}{\Delta t}\right) : \left(\frac{\widetilde{\mathbf{G}^p} - \mathbf{G}^p}{\Delta t}\right)} \tag{6.20}$$

where Δt is the time interval of the RK sub-step. Specific strain hardening models employed in this work will be described in the following sections along with initial conditions and material properties of the problem.

Because the plastic relaxation time scale τ_0 may not be known in many cases, we focus in this work on the limit of $\tau_0 \to 0$, where plastic relaxation to the yield surface occurs instantaneously at every RK sub-step. There are certain types of dispersed elastic-plastic waves which arise when the timescale τ_0 is comparable to the flow timescales [132], which this treatment will not resolve accurately in the vicinity of the shock. This assumption will replace a smooth plastic relaxation region with shocks, which are the "outer solution" based on the jump conditions in the governing equations. However, the numerical method described here is perfectly capable to simulate problems involving a finite τ_0, which will require less regularization, because the solution is smoother. Hence, the assumption of $\tau_0 \to 0$ is a worst case scenario for the numerical method from a regularization and robustness point of view, and the following descriptions of shock treatments will be less active when the flow is smoother.

6.2.4 Kinematic equations used in this work for multi-material deformation

Of the many possible formulations which can be used to compute elastic stresses and strain hardening, we select the $\underline{\mathbf{g}}^e$-$\underline{\mathbf{G}}^p$ formulation in this work for computational efficiency, and because it was found to be more stable than the alternatives in a previous study [5]. This choice saves 3 equations per material by using the symmetric tensor $\underline{\mathbf{G}}^p$ instead of $\underline{\mathbf{g}}^p$. Since $\underline{\mathbf{G}}^e$ and $\underline{\mathbf{G}}^p$ tensors are used directly to compute elastic stresses and equivalent plastic strain, this choice is also more convenient than the forms involving $\underline{\mathbf{g}}$ or $\underline{\mathbf{G}}$. Below we present the full equations as they are solved in this work, including numerical regularization terms on the right hand side of the transport equations. The subscript m represents each material

$$\underbrace{\frac{\partial \underline{\mathbf{g}}^e_m}{\partial t}}_{\substack{\text{local} \\ \text{derivative}}} + \underbrace{\nabla\left(\underline{\mathbf{g}}^e_m \cdot \boldsymbol{u}\right)}_{\substack{\text{curl-free} \\ \text{advection/strain}}} + \underbrace{\boldsymbol{u} \cdot \nabla\underline{\mathbf{g}}^e_m - \left(\nabla\underline{\mathbf{g}}^e_m\right)\cdot\boldsymbol{u}}_{\substack{\text{non-zero curl} \\ \text{advection/strain}}} - \underbrace{\mathbf{L}^p_m\underline{\mathbf{g}}^e_m}_{\text{elastic-plastic source}} = \underbrace{\frac{\zeta}{\Delta t}\overline{\left(\frac{\rho}{\rho_0\left|\underline{\mathbf{g}}^e_m\right|} - 1\right)}\underline{\mathbf{g}}^e_m}_{\text{determinant compatibility}} + \underbrace{\nabla\cdot\left(g^{e*}\nabla\underline{\mathbf{g}}^e_m\right)}_{\text{artificial diffusion}}$$

$$\tag{6.21}$$

$$\underbrace{\frac{\partial \underline{\mathbf{G}}_m^p}{\partial t}}_{\substack{\text{local} \\ \text{derivative}}} + \underbrace{\boldsymbol{u} \cdot \nabla \underline{\mathbf{G}}_m^p}_{\text{advection}} + \underbrace{\underline{\mathbf{L}}_m^p{}^T \underline{\mathbf{G}}_m^p + \underline{\mathbf{G}}_m^p \underline{\mathbf{L}}_m^p}_{\text{elastic-plastic source}} = \underbrace{\frac{\zeta}{\Delta t} \overline{\left(\frac{1}{|\underline{\mathbf{G}}_m^p|} - 1 \right) \underline{\mathbf{G}}_m^p}}_{\text{determinant compatibility}} + \underbrace{\nabla \cdot (g^{p*} \nabla \underline{\mathbf{G}}_m^p)}_{\text{artificial diffusion}} \tag{6.22}$$

where

$$\boldsymbol{L}_m^p = \frac{1}{2\mu_m \tau_{rel,m}} \underline{\mathbf{g}}_m^e \boldsymbol{\sigma}_m' \underline{\mathbf{g}}_m^{e\,-1} \tag{6.23}$$

$$\frac{1}{\tau_{\text{rel,m}}} = \frac{1}{(\rho_m/\rho_{0,m})\,\tau_{0,m}} \left[\frac{\mathrm{R}\left(\boldsymbol{\sigma}_m' : \boldsymbol{\sigma}_m' - \frac{2}{3}\sigma_{Y,m}^2 \right)}{\mu_m^2} \right], \tag{6.24}$$

The left-hand side of the kinematic equations contain terms which are physical: spatially local time derivative, advection, and strain, as well as the modeled elastic-plastic source terms (discussed in Section 6.2.3) which capture the effect of plastic yielding. The right-hand-side contains the numerical regularization terms, including the determinant compatibility correction and artificial diffusion (discussed in Section 6.2.5). The determinant compatibility correction originates from Miller and Colella [135], in which Δt is the numerical time step, and ζ controls the magnitude of the correction and is limited by the stability of the method. $\zeta = 1/6$ is used here. The overbar over the determinant compatibility terms represents a truncated Gaussian filter applied along each grid direction (details in Cook [39]). This filter is applied to prevent coupling with the localized artificial diffusivity, which responds to high wavenumber noise in the solution variables. This filter also reduces the magnitude of the correction, effectively imposing this constraint more gently.

Time-advancing kinematic equations for each material can be expensive, as this leads to a system of $4+16M$ PDEs, where M is the number of species. This number can be reduced by instead solving equations for mixture $\geq$ and mixture $\underline{\mathbf{G}}^p$, using the method described in Subramaniam [187] Chapter 8, where it is applied to the $\underline{\mathbf{G}}^e$ equations. This reduced formulation is used in this work only for the Taylor impact test problem in section 6.4.

6.2.5 Localized artificial diffusivity

The localized artificial diffusivity framework used here is based largely on Cook [39], with some modifications from Kawai and Lele [108] and Ghaisas et al. [73]. The localized artificial diffusivities which have been used previously are the artificial shear viscosity, μ^*; artificial bulk viscosity, β^*; artificial thermal conductivity, κ^*; and artificial diffusivity of species m, D_m^*. These are described below. In these expressions, the overbar denotes a truncated Gaussian filter applied along each grid direction (details in Cook [39]); Δ_i is the grid spacing in the i direction; $\Delta_{i,\mu}$, $\Delta_{i,\beta}$, $\Delta_{i,\kappa}$, and Δ_{i,Y_m} are weighted grid length scales in direction i; c_s is the linear longitudinal wave (sound) speed; H is

the Heaviside function; and $\varepsilon = 10^{-32}$ [1].

$$\mu^* = C_\mu \rho \overline{\left| \sum_{k=1}^{3} \frac{\partial^r S}{\partial x_k^r} \Delta_k^r \Delta_{k,\mu}^2 \right|}; \qquad \Delta_{i,\mu} = \Delta_i. \tag{6.25}$$

$$\beta^* = C_\beta \rho \overline{\left| \sum_{k=1}^{3} \frac{\partial^r (\boldsymbol{\nabla} \cdot \boldsymbol{u})}{\partial x_k^r} \Delta_k^r \Delta_{k,\beta}^2 \right|}; \qquad \Delta_{i,\beta} = \Delta_i \frac{\left(\frac{\partial \rho}{\partial x_i} \right)^2}{\sum_{k=1}^{3} \left(\frac{\partial \rho}{\partial x_k} \right)^2 + \varepsilon}. \tag{6.26}$$

$$\kappa^* = C_\kappa \frac{\rho c_s}{T} \overline{\left| \sum_{k=1}^{3} \frac{\partial^r e_h}{\partial x_k^r} \Delta_k^r \Delta_{k,\kappa} \right|}; \qquad \Delta_{i,\kappa} = \Delta_i \frac{\left(\frac{\partial e_h}{\partial x_i} \right)^2}{\sum_{k=1}^{3} \left(\frac{\partial e_h}{\partial x_k} \right)^2 + \varepsilon}. \tag{6.27}$$

$$D_m^* = \max \left\{ C_D c_s \overline{\left| \sum_{k=1}^{3} \frac{\partial^r Y_m}{\partial x_k^r} \Delta_k^r \Delta_{k,D} \right|}, \; C_Y \frac{c_s}{2} \overline{\left(|Y_m| - 1 + |1 - Y_m| \right) \sum_{k=1}^{3} \Delta_{k,Y}} \right\};$$

$$\Delta_{i,D} = \Delta_i \frac{\left(\frac{\partial Y_m}{\partial x_i} \right)^2}{\sum_{k=1}^{3} \left(\frac{\partial Y_m}{\partial x_k} \right)^2 + \varepsilon}; \qquad \Delta_{i,Y} = \Delta_i \frac{\left| \frac{\partial Y_m}{\partial x_i} \right|}{\sqrt{\sum_{k=1}^{3} \left(\frac{\partial Y_m}{\partial x_k} \right)^2 + \varepsilon}}. \tag{6.28}$$

Here, $S = \sqrt{\boldsymbol{S} : \boldsymbol{S}}$ is a norm of the strain rate tensor. These artificial diffusivities are the same as those suggested by Kawai and Lele [108], Mani et al. [122], and Kawai et al. [107] have explored variations on the artificial bulk viscosity which further localize it with a Ducros-like switching function and a Heaviside function to restrict it to compression. Here, we do not employ a switching function in order to effectively smooth tensile shocks, which can occur in solids, unlike in gasdynamics.

To regularize large jumps in strain, new artificial diffusivities are introduced for each kinematic tensor: g^{e*} and g^{p*}. These are calcualted for each material, but the subscript "m" is dropped for clarity.

$$g^{e*} = C_{ge} c_s (1 - f_{sw}) \overline{\left| \sum_{k=1}^{3} \frac{\partial^r \epsilon_e}{\partial x_k^r} \Delta_k^r \Delta_{k,ge} \right|}; \qquad \Delta_{i,ge} = \Delta_i \frac{\left(\frac{\partial \epsilon_e}{\partial x_i} \right)^2}{\sum_{k=1}^{3} \left(\frac{\partial \epsilon_e}{\partial x_k} \right)^2 + \varepsilon}. \tag{6.29}$$

$$g^{p*} = C_{gp} c_s (1 - f_{sw}) \overline{\left| \sum_{k=1}^{3} \frac{\partial^r \epsilon_p}{\partial x_k^r} \Delta_k^r \Delta_{k,gp} \right|}; \qquad \Delta_{i,gp} = \Delta_i \frac{\left(\frac{\partial \epsilon_p}{\partial x_i} \right)^2}{\sum_{k=1}^{3} \left(\frac{\partial \epsilon_p}{\partial x_k} \right)^2 + \varepsilon}. \tag{6.30}$$

---

[1] We note that wherever ε appears, it must be multiplied by a factor with appropriate dimensions to maintain dimensional consistency and perform its intended expected function as the magnitudes of flow variables vary across simulation cases. Appropriate dimensional constants can be constructed from positive quantities, such as the speed of sound, minimum grid spacing, density, and temperature.

$$f_{sw} = \frac{(\boldsymbol{\nabla} \cdot \boldsymbol{u})^2}{(\boldsymbol{\nabla} \cdot \boldsymbol{u})^2 + |\boldsymbol{\nabla} \times \boldsymbol{u}|^2 + \varepsilon}. \tag{6.31}$$

Here, $\epsilon_e = \sqrt{\frac{2}{3}\boldsymbol{e}_e : \boldsymbol{e}_e}$ is a norm of the Euler-Almansi finite strain tensor associated with elastic deformation, where $\boldsymbol{e}_e = \frac{1}{2}(\mathbf{1} - \underline{\mathbf{G}}^e)$. The plastic strain norm ϵ_p is defined analogously: $\epsilon_p = \sqrt{\frac{2}{3}\boldsymbol{e}_p : \boldsymbol{e}_p}$, where $\boldsymbol{e}_p = \frac{1}{2}(\mathbf{1} - \underline{\mathbf{G}}^p)$. The factor of $2/3$ in these norms is included so that the plastic strain norm is the "equivalent plastic strain", which takes into account the assumptions of isochoric plastic flow.

We observe that LAD is not strictly necessary to ensure stability for the $\underline{\mathbf{g}}^e$ equations; in fact, it has not been included in previous formulations [73, 188], because the elastic deformation is typically small in magnitude for metals. However, LAD is necessary to provide stability for the $\underline{\mathbf{G}}^p$ equations, especially in problems involving significant shear deformation. Accordingly, the switching function $(1 - f_{sw})$ localizes artificial diffusion of kinematic tensors to regions dominated by vorticity. This switching function and notation is inspired by the Ducros-like switching function introduced by Kawai et al. [107], but we use $(1 - f_{sw})$ to obtain the opposite effect. It was found that artificial bulk viscosity provided sufficient regularization at jumps in normal strain, so this switching function ensures that superfluous dissipation is not added.

For problems involving material strength, the following coefficient values were found to be effective (with more details on provided in section 6.3): Typical values for the model coefficients are $C_\mu = 0.002$, $C_\beta = 5$, $C_\kappa = 1 \times 10^{-2}$, $C_D = 3 \times 10^{-3}$, $C_Y = 1 \times 10^2$, $C_{ge} = 1$, and $C_{gp} = 1 \times 10^{-2}$.

6.2.6 Treatment of rotation

Another important ingredient of the numerical method is special treatment given to rotation. Ortega et al. [148] observed that resolving rotation in a multi-material simulation can lead to physical problems when two initially disconnected regions of material with different rotation histories reconnect. This can occur in the Richtmyer-Meshkov instability in converging geometries, which they studied in later work [149]. They resolved this problem by discarding the information about rotation, since it does not contribute to elastic energy or elastic stresses. This is done by considering the polar decomposition of the elastic inverse deformation gradient $\underline{\mathbf{g}}^e = \underline{\mathbf{R}}^e \underline{\mathbf{v}}^e$, where $\underline{\mathbf{R}}^e$ is a rotation tensor and $\underline{\mathbf{v}}^e$ is a symmetric strain tensor. Their approach amounts to solving an equation for $\underline{\mathbf{g}}^e$, but at the end of each timestep, re-assigning it to be symmetric:

$$\underline{\mathbf{g}}^e_{\,n+1} \leftarrow \underline{\mathbf{v}}^e_{\,n+1} \tag{6.32}$$

We follow the same approach here, as we find that it is also beneficial for properly resolving elastic stresses and energy under conditions of severe shear. Because we are solving an equation for the plastic Finger tensor, the rotation is moved to $\underline{\mathbf{G}}^p$ rather than discarded, in order to preserve the

total deformation ($\underline{\mathbf{G}} = \underline{\mathbf{g}}^{eT}\underline{\mathbf{G}}^{p}\underline{\mathbf{g}}^{e}$)

$$\underline{\mathbf{G}}^{p}_{n+1} \leftarrow \underline{\mathbf{R}}^{eT}_{n+1}\underline{\mathbf{G}}^{p}_{n+1}\underline{\mathbf{R}}^{e}_{n+1}. \tag{6.33}$$

This treatment of rotation does not alter the invariants of either $\underline{\mathbf{g}}^{e}$ or $\underline{\mathbf{G}}^{p}$. Consequently, it does not affect the elastic energy or equivalent plastic strain. Because the elastic stresses are computed based on the elastic Finger tensor ($\underline{\mathbf{G}}^{e} = \underline{\mathbf{v}}^{e}\underline{\mathbf{v}}^{e}$), the elastic stresses are also unaffected by this procedure.

6.2.7 Equation of state and constitutive equation

The above system of PDEs is closed with an equation of state (EOS). In this work, we consider only hyperelastic materials, in which elastic stresses are derived from an energy functional for thermodynamic consistency. It is assumed that the internal energy is additively decomposed into hydrodynamic and elastic components. The hydrodynamic part (e_h) depends only on pressure and density. The elastic energy (e_e) depends only on invariants of the elastic Finger tensor:

$$e = e_h\left(p, \rho\right) + e_e\left(\hat{\boldsymbol{g}}\right), \tag{6.34}$$

in which $\hat{\boldsymbol{g}} = |\underline{\boldsymbol{G^e}}|^{-1/3}\underline{\boldsymbol{G^e}}$.

This split allows for the use of wide variety of equations of state, but for purposes of demonstration in this work, we primarily represent materials as stiffened gases according to

$$e_h = \frac{p + \gamma p_\infty}{\left(\gamma - 1\right)\rho}, \tag{6.35}$$

where p_∞ (with units of pressure) and γ (non-dimensional) are material constants of the stiffened gas model for the hydrodynamic component of internal energy. For the elastic energy we use the relationship from Favrie and Gavrilyuk [59]

$$e_e = \frac{\mu}{4\rho_0}\text{tr}\left(\left(\hat{\boldsymbol{g}} - \underline{\mathbf{1}}\right)^2\right), \tag{6.36}$$

where μ is the material shear modulus. With this EOS, the Cauchy stress, $\underline{\boldsymbol{\sigma}}$, satisfying the Clausius-Duhem inequality is

$$\underline{\boldsymbol{\sigma}} = -p\underline{\boldsymbol{I}} - \mu\frac{\rho}{\rho_0}\left(|\underline{\boldsymbol{G^e}}|^{-2/3}\,\text{dev}\left(\left(\underline{\boldsymbol{G^e}}\right)^2\right) - |\underline{\boldsymbol{G^e}}|^{-1/3}\,\text{dev}\left(\underline{\boldsymbol{G^e}}\right)\right), \tag{6.37}$$

in which $\text{dev}\left(\underline{\boldsymbol{G^e}}\right)$ signifies the deviatoric component of the tensor: $\text{dev}\left(\underline{\boldsymbol{G^e}}\right) = \underline{\boldsymbol{G^e}} - \text{tr}\left(\underline{\boldsymbol{G^e}}\right)/3$, with $\text{tr}\left(\underline{\boldsymbol{G^e}}\right)$ signifying the trace of the tensor. The elastic component of the internal energy, e_e, is assumed to be isentropic. Therefore, the temperature, T, and entropy, η, are defined by the

hydrodynamic stiffened gas component of the EOS, as follows.

$$e_h = C_v T \left(\frac{p + \gamma p_\infty}{p + p_\infty} \right) \qquad R = C_p - C_v \qquad \gamma = \frac{C_p}{C_v}$$

$$\eta - \eta_0 = C_p \ln \left(\frac{T}{T_0} \right) + R \ln \left(\frac{p_0 + p_\infty}{p + p_\infty} \right). \tag{6.38}$$

Here, η_0 is the reference entropy at pressure, p_0, and temperature T_0.

It is important to emphasize that the numerical method described here is not limited to stiffened gases and this particular choice of elastic energy functional, provided the equation of state and energy functional can be represented according to equation 6.34, and that the functional form of Cauchy stress is appropriately derived. This permits more realistic EOS, such as Mie-Grüneisen to be used without issue, as is done in Section 6.4 on the Taylor impact problem.

To close the system of PDEs for multi-material mixtures we assume pressure-temperature equilibrium among the phases, and solve for the pressure and temperature which satisfy the mixture density and energy according to the method in Cook [40], Appendix B, which is valid for any equation of state, including tabular.

6.2.8 High-order numerical method

The numerical discretization follows the method previously documented in Cook [39], Ghaisas et al. [73], and Subramaniam et al. [188]. The equations are discretized on an Eulerian, uniform, Cartesian grid. Spatial derivatives are computed using a high-resolution, penta-diagonal, 10th-order, compact finite difference scheme. This scheme is applied in the domain interior and near the boundaries in the cases of symmetry, anti-symmetry or periodic boundary conditions. Otherwise boundary derivatives are reduced to a 4th-order, one-sided, compact difference scheme. Time advancement is achieved using a five-stage, 4th-order, Runge-Kutta method from Kennedy et al. [109], with an adaptive time step based on a CFL condition.

The CFL time step limit in 3-D is given by

$$\Delta t_{CFL} = MIN \left[\left(\frac{|u|}{\Delta_x} + \frac{|v|}{\Delta_y} + \frac{|w|}{\Delta_z} + c_s \sqrt{\frac{1}{\Delta_x^2} + \frac{1}{\Delta_y^2} + \frac{1}{\Delta_z^2}} \right)^{-1} \right] \tag{6.39}$$

, where $MIN[\,]$ represents the minimum taken over the whole domain, u, v, and w are the Cartesian velocity components, Δ_x, Δ_y, and Δ_z, are the Cartesian grid spacings, and c_s is the longitudinal wave speed, which for the materials considered here is

$$c_s = \sqrt{\frac{\gamma(p + p_\infty)}{\rho} + \frac{4}{3} \frac{\mu}{\rho}}. \tag{6.40}$$

The time step limits for the artificial diffusivity terms are given by

$$\Delta t_\mu = MIN \left[\frac{\rho \Delta^2}{\mu^*} \right] \tag{6.41}$$

$$\Delta t_\beta = MIN \left[\frac{\rho \Delta^2}{\beta^*} \right] \tag{6.42}$$

$$\Delta t_\kappa = MIN \left[\left(\frac{\rho \Delta^4}{\kappa^* T} \right)^{1/3} \right] \tag{6.43}$$

$$\Delta t_D = MIN \left[\frac{\Delta^2}{D_m{}^*} \right] \tag{6.44}$$

$$\Delta t_{ge} = MIN \left[\frac{\Delta^2}{g^{e*}} \right] \tag{6.45}$$

$$\Delta t_{gp} = MIN \left[\frac{\Delta^2}{g^{p*}} \right] \tag{6.46}$$

$$\tag{6.47}$$

where $\Delta = min(\Delta_x, \Delta_y, \Delta_z)$. At each simulation time step, the time step is calculated according to

$$\Delta t = min \left(\text{CFL} \Delta t_{\text{CFL}}, \ 0.2 \Delta t_\beta, \ 0.2 \Delta t_\kappa, \ 0.2 \Delta t_D, \ 0.2 \Delta t_\mu, \ 0.2 \Delta t_{ge}, \ 0.2 \Delta t_{gp} \right) \tag{6.48}$$

where CFL ≤ 1 is the CFL number used for the simulation, and 0.2 is used to keep the time step within the various diffusive stability limits. A high-resolution, penta-diagonal, 8th-order, compact Padé filter, with cutoff parameters described by Ghaisas et al. [73] is applied after each stage of the Runge-Kutta algorithm to each of the conserved variables and kinematic tensors – ρY_m, $\rho \boldsymbol{u}$, $\rho \left(e + \frac{1}{2} \boldsymbol{u} \cdot \boldsymbol{u} \right)$, $\underline{\mathbf{g}}^e_m$, and $\underline{\mathbf{G}}^p_m$. This filter removes approximately the top 10% of the grid-resolvable wavenumber content, thereby mitigating aliasing errors and numerical instability in the high-wavenumber range, and preventing the artificial properties from becoming too large [39].

6.2.9   Damping of kinematic equations in fluid regions

For multi-phase problems involving solid-fluid interfaces, the lack of material strength in the fluid regions can lead to very large deformations, which do not contribute meaningfully to elastic stresses or elastic energy, and which can cause the equations for $\underline{\mathbf{g}}^e$ and $\underline{\mathbf{G}}^p$ to blow up. To remedy this, the kinematic tensors are damped when the solid volume fraction is below a threshold. In these regions, the kinematic tensors are relaxed back toward the identity tensor by overwriting them at the end of

each RK step according to

$$\underline{\mathbf{g}}^e \leftarrow \underline{\mathbf{g}}^e + \xi \mathbf{1} \tag{6.49}$$

$$\underline{\mathbf{G}}^p \leftarrow \underline{\mathbf{G}}^p + \xi \mathbf{1} \tag{6.50}$$

$$\xi = 1 - \min\left(\max\left(\frac{\phi_s}{\phi_{thresh}}, 0\right), 1\right) \tag{6.51}$$

where ϕ_s is the volume fraction of the solid(s), and ϕ_{thresh} is a user-specified threshold, here taken to be 0.01.

6.3 Numerical Tests

The following suite of numerical tests is provided as verification and validation of the numerical method. These tests demonstrate the convergence of the method, how the localized artificial diffusivity captures shocks in elastic-plastic solids, expected physical behavior with and without strain hardening in impact problems, and multi-scale resolution properties. All tests in this section involve only a single material.

6.3.1 Order of convergence

To assess the order of convergence in time and space, a problem involving weak elastic waves is considered. All model terms (plasticity, LAD, determinant compatibility) are omitted to demonstrate the convergence behavior in space and time based on the properties of the spatial differencing and time integration schemes used. In the problem considered, the boundary conditions are periodic, and the following initial condition is used:

$$u^*(x^*, t_0) = 0.0022 \sin(2\pi x^*) \tag{6.52}$$

$$v^*(x^*, t_0) = 0.0022 \sin(2\pi x^*) \tag{6.53}$$

$$p^*(x^*, t_0) = 0 \tag{6.54}$$

$$\rho^*(x, t_0) = 1 \tag{6.55}$$

$$\underline{\mathbf{g}}^e(x, t_0) = \underline{\boldsymbol{\delta}}. \tag{6.56}$$

This problem has been non-dimensionalized by the domain length L_x, the undeformed density ρ_0, and the pressure scale $p_{ref} = \gamma p_\infty + 4/3\mu$, which results in a velocity scale equal to the speed of sound, c_s. This initial condition results in small-amplitude elastic shear and normal deformation. For both spatial and temporal convergence, the initial value problem is allowed to evolve until $t^* = 0.13$. The initial and final profiles of u, v, and ϵ_e obtained with the spatially refined solution are plotted in Figure 6.1. The initial velocity profiles induce elastic stresses, which decrease the magnitude of
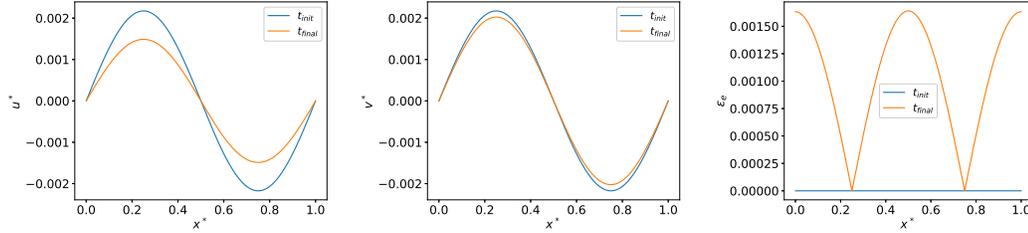
FIGURE 6.1: Reference solution obtained with $N_x = 2000$ for the elastic convergence test. Initial and final profiles are plotted for u, v, and ϵ_e.

the velocities.

Convergence is assessed by computing the L_2 norm of the error at 10 sampling points ($x^* = 0, 0.1, 0.2, ..., 0.9$) with respect to a highly refined reference solution. For spatial convergence, a constant timestep of $\Delta t^* = 10^{-5}$ is used as N_x is varied from 10 to 10^3, to ensure that the errors due to temporal integration are kept small. The refined solution uses $N_x = 10^4$. For temporal convergence, a constant grid of $N_x = 100$ is used as Δt^* is varied from 1.3×10^{-2} to 5×10^{-4}. For the refined temporal solution, a finer grid ($N_x = 6000$) is used, with $\Delta t^* = 2e - 4$. The convergence properties in time and space are shown in Figure 6.2. This is done without the rotation treatment in Figures 6.2(a,b), and with it in in Figures 6.2(c,d). Regardless, 8th order spatial convergence is obtained. Although 10th order accurate compact differencing schemes are used, the observed convergence behavior is expected due to the 8th order compact filter which is applied for stability purposes. Without the rotation treatment, 4th order temporal convergence is obtained, but this is reduced to 1st order when the rotation treatment is applied. In a purely elastic problem ,the rotation treatment is not necessary, but in plastic problems, it is needed to prevent excessive rotation of material elements which become unresolvable and lead to instability. Because the implicit scheme used to relax back to the yield surface during plastic deformation is first order in time [59], the first order temporal convergence obtained from the rotation treatment is not a concern in problems of interest, which involve significant plastic deformation.

6.3.2 Shock capturing

In elastic-plastic solids, the familiar discontinuities of gasdynamics – shocks, contact surfaces, and material interfaces – still occur, but the shocks come in a wider variety. Shocks may be purely elastic or elastic-plastic, and they may involve shear deformation, since solids generate stress due to shear strain. These additional discontinuities are treated here with localized artificial diffusivity. In this section, we demonstrate how normal and shear strain discontinuities are regularized. This is accomplished with the artificial bulk viscosity β^*, plus the new artificial diffusivities, g^{e*} and g^{p*}. Here we present how to tune the coefficients for these terms: C_β, C_{ge}, and C_{gp}. The coefficients
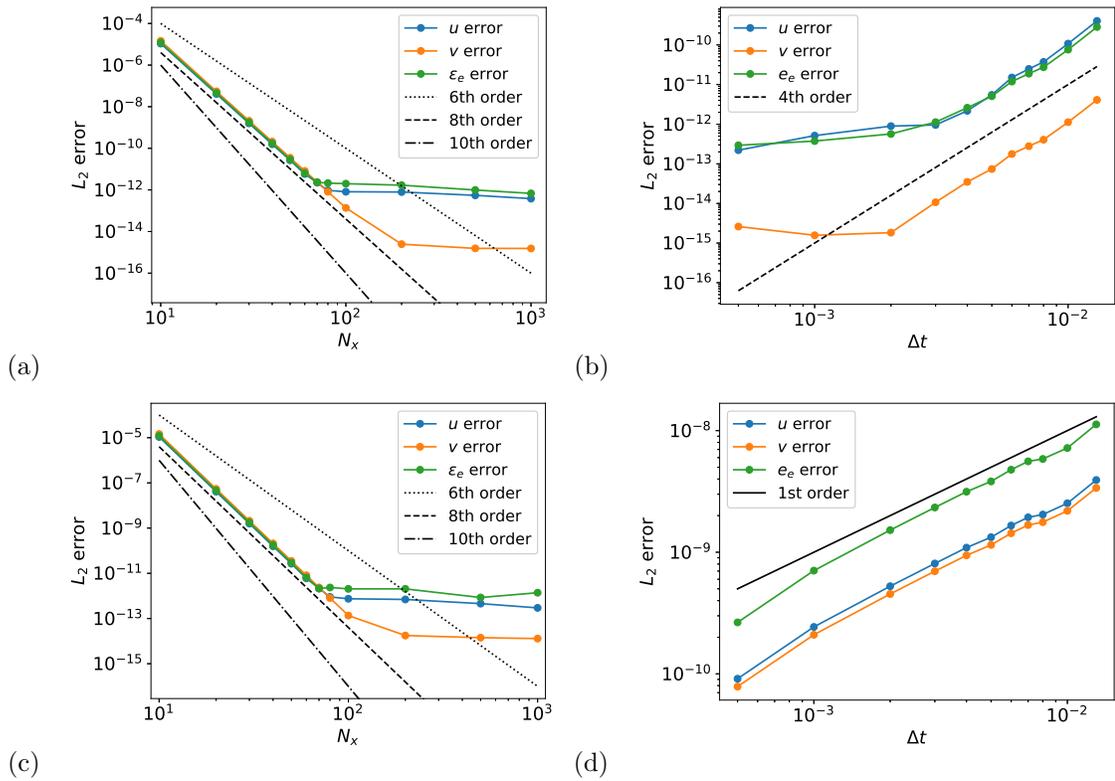
FIGURE 6.2: Convergence behavior for an elastic wave propagation problem. (a) Spatial convergence (8th order) without the rotation treatment applied. (b) Temporal convergence (4th order) without the rotation treatment applied. (c) Spatial convergence (8th order) with the rotation treatment applied. (d) Temporal convergence (1st order) with the rotation treatment applied.

for artificial shear viscosity, thermal conductivity, and mass diffusivity are left unchanged from the previous recommendations of Cook [39]: $C_\mu = 0.002$, $C_\kappa = 0.01$, $C_D = 0.003$, and $C_Y = 100$.

The approach taken to tune the coefficients C_β and C_{ge} is to use 1D simulations of single material in a moving frame of reference such that the shock is stationary in the domain. The simulations are initialized with a sharp discontinuity in the middle of the domain, with the states on either side of the shock based on the appropriate jump conditions, which are discussed in the following sections. The simulation is run until a wave traveling at the speed of sound can traverse the domain 50 times, so that the LAD terms and discontinuity reach a steady state. Then, the shock thickness and Gibbs phenomena are assessed. Across all tests, a grid of $N_x = 101$ is used, with $CFL = 0.5$. The solution at the boundaries is Gaussian filtered several times to approximate a non-reflecting boundary condition. For the purposes of the following discussion, all discontinuities are considered as 1D jumps along the x axis. We use "normal" to mean the x direction and u velocity, and "transverse" to mean y and v. Without loss of generality, we do not consider shear deformation in z, and therefore $w = 0$ throughout this section.

Normal shocks

The shock-jump relations for normal shocks in elastic-plastic solids are

$$\rho_1 u_1 = \rho_2 u_2 \qquad (6.57)$$

$$\rho_1 u_1^2 - \sigma_{11}^{(1)} = \rho_2 u_2^2 - \sigma_{11}^{(2)} \qquad (6.58)$$

$$\rho_1 u_1 (\epsilon_1 + \frac{1}{2} u_1^2 + \frac{1}{2} v_1^2) - \sigma_{11}^{(1)} u_1 - \sigma_{21}^{(1)} v_1 = \rho_2 u_2 (\epsilon_2 + \frac{1}{2} u_2^2 + \frac{1}{2} v_2^2) - \sigma_{11}^{(2)} u_2 - \sigma_{21}^{(2)} v_2, \qquad (6.59)$$

which represent conservation of mass, momentum, and energy, respectively. The subscripts "1" and "2" or superscripts (2) and (2) denote states on either side of the shock. In normal shocks, the energy equation can be further simplified because the terms involving v and σ_{21} are zero.

To solve for the states on either side of a purely elastic normal shock, the approach is similar to that used for gasdynamics, with a few additions. In this case, $\underline{\mathbf{g}}^e = \underline{\mathbf{g}}$ and $\underline{\mathbf{g}}^p = \mathbf{1}$. The determinant compatibility condition is used to express $g_{11} = \rho/\rho_0$. Adding to the jump relations the equation of state and constitutive equation results in a system of 5 equations with 6 unknowns. Specifying a pressure ratio p_2/p_1 or an impact velocity $\Delta U = u_2 - u_1$ fully constrains the system, which can be solved numerically.

We consider purely elastic normal shocks for three hypothetical materials, using a copper-like stiffened gas as a baseline. The material properties of this material are $p_\infty = 68.23\ GPa$, $\gamma = 2$, $\rho_0 = 8930\ kg/m^3$, $\mu = 39.38\ GPa$, and $\sigma_Y = 0.12\ GPa$. To test the dependence of the recommended LAD coefficients on material properties, we also consider the same material with the yield stress increased and decreased by a factor of 10 ($0.1\sigma_{cu}$ and $10\sigma_{cu}$, respectively). For each material, we
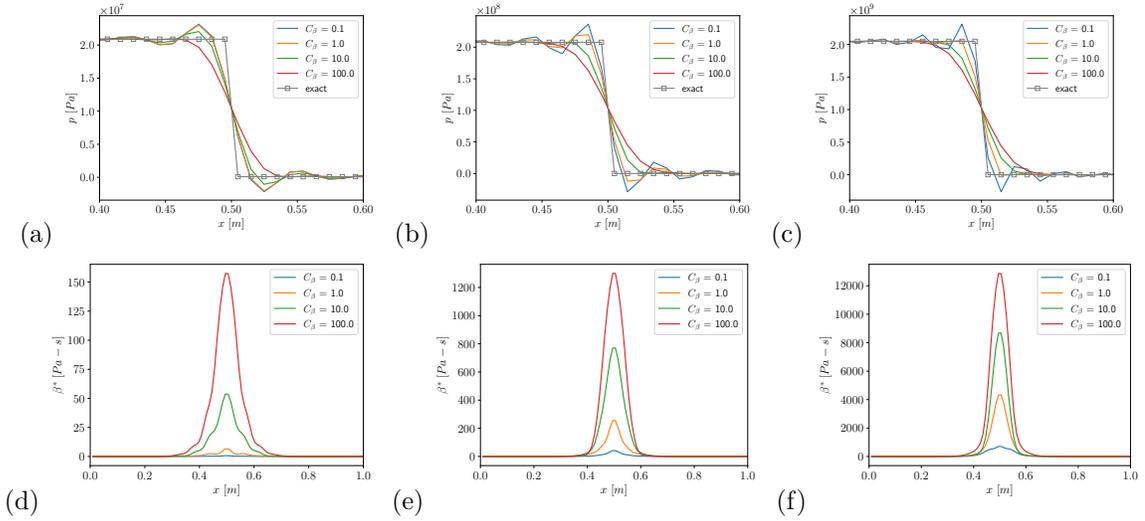
FIGURE 6.3: Regularization of elastic normal shocks as C_β is varied. The sub-figures show the effect on (a) pressure for $\sigma_Y = 0.1\sigma_{cu}$, (b) pressure for $\sigma_Y = \sigma_{cu}$, (c) pressure for $\sigma_Y = 10\sigma_{cu}$, (d) artificial bulk viscosity for $\sigma_Y = 0.1\sigma_{cu}$, (e) artificial bulk viscosity for $\sigma_Y = \sigma_{cu}$, and (f) artificial bulk viscosity for $\sigma_Y = 10\sigma_{cu}$.

consider the elastic precursor, which is the strongest purely elastic shock a material will support, which amounts to changing the shock strength. In this test, the code is run in purely elastic mode. The behavior of the LAD when plasticity is turned on is captured in a later test.

Figure 6.3(a-c) shows how varying C_β can be used to control shock spreading in the 3 materials considered. In all cases, the artificial bulk viscosity β^*, shown in Figure 6.3(d-f) is localized to the region around the shock, and it increases in magnitude as C_β and the shock strength are increased. As C_β is increased, the shocks become thicker and the Gibbs phenomena around the jumps decrease in amplitude. This tradeoff is quantified in Figure 6.4. Wiggle amplitude is calculated as the height of the Gibbs phenomena over the plateau divided by the size of the jump. The shock width is defined the length scale associated with the slope at the shock midpoint and the jump magnitude. A value of C_β can be selected by identifying a threshold for acceptable amplitude of wiggles and choosing the smallest C_β which achieves that. The ideal value of C_β does depend on the material, because of the difference in shock strength. For the metal-like materials considered here a value of $C_\beta = 5$ is recommended to balance shock thickness and Gibbs phenomena amplitude.

If a shock is strong enough to cause yielding, the procedure for obtaining the shock states is more complex. In this situation, two shocks form: an elastic precursor, which represents the strongest possible elastic shock before the material yields (jump 1-2), and behind it a slower-moving elastic-plastic shock, in which the rest of the deformation specified by the initial condition occurs (jump 3-4) [132]. The first step is to solve for the elastic precursor, which uses the same set of 5 equations as a purely elastic shock. However, the shock strength must be determined by finding the weakest
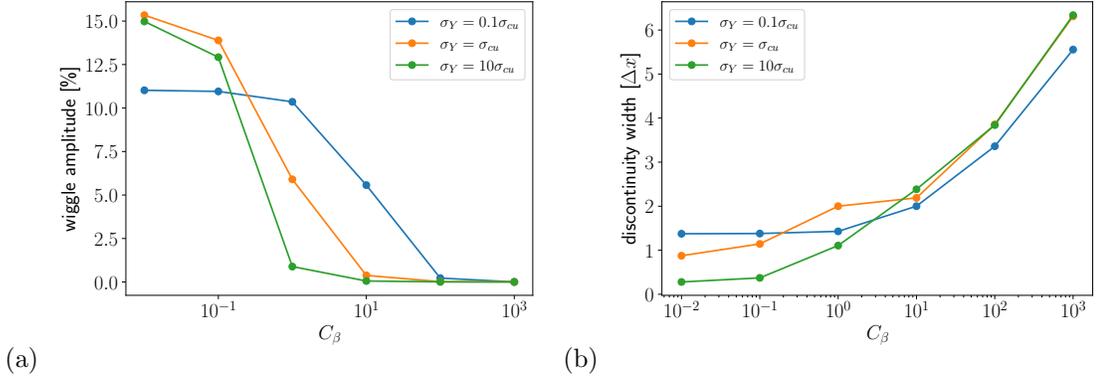
FIGURE 6.4: Effect of C_β in elastic normal precursor shocks in 3 different materials. (a) Amplitude of Gibbs phenomena around the shock. (b) Width of the shock.

shock which results in a stress on the yield surface. This adds the yield condition $\underline{\boldsymbol{\sigma}}' : \underline{\boldsymbol{\sigma}}' = \frac{2}{3}\sigma_Y^2$, so now the system of 6 equations and 6 unknowns is closed and can be solved numerically.

To solve for the states on either side of the elastic-plastic shock, the same 6 equations and 6 unknowns can be used, but to these we add the requirement that the velocity jump be equal to the overall velocity minus the jump associated with the precursor (+1 equation). The internal variables $(\rho, e, \underline{\mathbf{g}}^e, p)$ of state 3 are the same as state 2. The lo of freedom needed for state 4 needed comes from g_{11}^e, which no longer equals ρ/ρ_0. Instead, $\underline{\mathbf{g}}^e$ and $\underline{\mathbf{g}}^p$ have the following structure

$$
\underline{\mathbf{g}} = \begin{pmatrix} \frac{\rho}{\rho_0} & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} = \underline{\mathbf{g}}^p \underline{\mathbf{g}}^e = \begin{pmatrix} \frac{\rho}{\rho_0 g_{11}^e} & 0 & 0 \\ 0 & \sqrt{\frac{\rho_0 g_{11}^e}{\rho}} & 0 \\ 0 & 0 & \sqrt{\frac{\rho_0 g_{11}^e}{\rho}} \end{pmatrix} \begin{pmatrix} g_{11}^e & 0 & 0 \\ 0 & \sqrt{\frac{\rho}{\rho_0 g_{11}^e}} & 0 \\ 0 & 0 & \sqrt{\frac{\rho}{\rho_0 g_{11}^e}} \end{pmatrix} \quad (6.60)
$$

in which we have made use of the fact that in normal deformation in an infinite medium, there should be no difference between the deformation states in y and z, and that the determinants of the kinematic tensors are related to the density: $\det \underline{\mathbf{g}} = \det \underline{\mathbf{g}}^p \det \underline{\mathbf{g}}^e = (1)(\rho/\rho_0)$.

To verify that β^* can also regularize elastic-plastic normal shocks, we consider the jump from the elastic precursor state (3) to the plastically deformed state (4). We consider a shock strength associated with an impact of $\Delta U = \pm 100 \ m/s$. The same three materials are considered, and the effect of varying C_β on profiles of pressure and artificial bulk viscosity is shown in Figure 6.5. For the material with lowest yield strength $- \sigma_Y = 0.1\sigma_{cu}$, shown in Figure 6.5(a,d) – the Gibbs phenomena are affected similarly to the elastic shock. For the materials with higher yield strength, shown in Figure 6.5(b,c,e,f), The Gibbs phenomena do not appear as severe, even with small values of C_β. This is due to the interaction of Gibbs phenomena with the implicit plastic relaxation. The Gibbs phenomena arise due to compact derivative and filter operators, and these fluctuations can cause
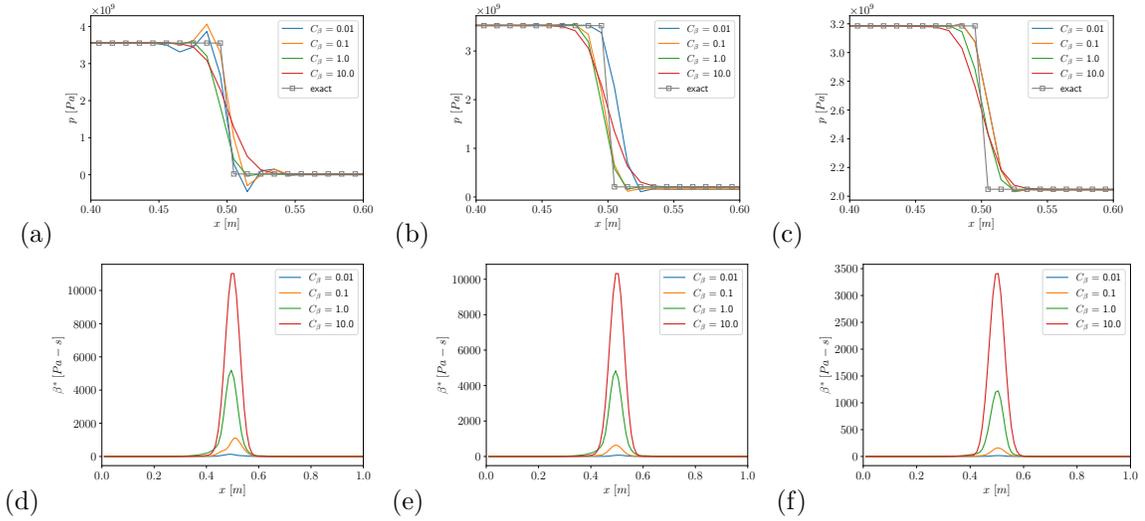
FIGURE 6.5: Regularization of elastic-plastic normal shocks as C_β is varied for an overall impact strength of $u = \pm 100 \ m/s$. The sub-figures show the effect on (a) pressure for $\sigma_Y = 0.1\sigma_{cu}$, (b) pressure for $\sigma_Y = \sigma_{cu}$, (c) pressure for $\sigma_Y = 10\sigma_{cu}$, (d) artificial bulk viscosity for $\sigma_Y = 0.1\sigma_{cu}$, (e) artificial bulk viscosity for $\sigma_Y = \sigma_{cu}$, and (f) artificial bulk viscosity for $\sigma_Y = 10\sigma_{cu}$.

stresses higher than yield, which are relaxed back to the yield surface at every RK sub-step. Thus, the implicit plastic relaxation provides some additional regularization around shocks causing plastic deformation.

The exercise from Figure 6.5 was repeated for weaker ($\Delta U = \pm 10 \ m/s$) and stronger ($\Delta U = \pm 1000 \ m/s$) shocks, and the effect of C_β on wiggle amplitude and shock width is summarized in Figure 6.6. For the weakest shock strength, the Gibbs phenomena are strongly attenuated by plastic relaxation, as seen in Figure 6.6(a). The highest yield strength material is omitted because this shock is too weak to cause it to yield. This Gibbs phenomena are also damped for the moderate shock strength, with the materials at higher yield strength in Figure 6.6(b). In both cases the shocks are weaker, and therefore the Gibbs phenomena cause smaller excesses over the yield criterion, which can be more effectively damped by the implicit relaxation. Again, a value of $C_\beta = 5$ is effective across the board without thickening shocks excessively.

Shear discontinuities

The shock-jump relations from normal shocks also apply to shear shocks, but they are augmented with two additional relations, which describe the jump in transverse momentum and the jump in kinematic variables. The jump condition for transverse momentum is

$$\rho_1 u_1 v_1 - \sigma_{21}^{(1)} = \rho_2 u_2 v_2 - \sigma_{21}^{(2)} \tag{6.61}$$
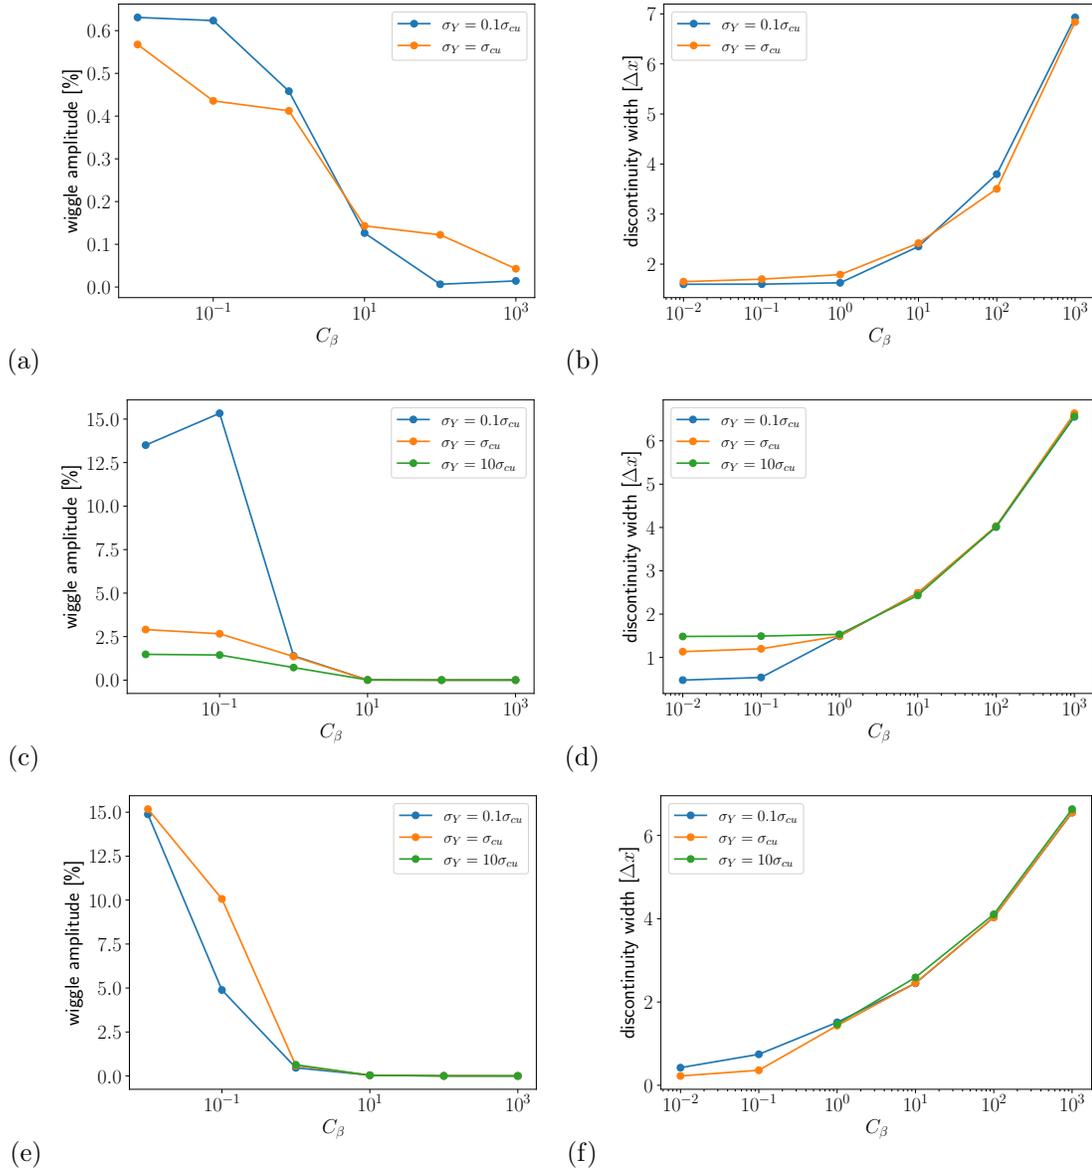
FIGURE 6.6: Effect of C_β in elastic-plastic normal shocks in 3 different materials and 3 different shock strengths. (a) Amplitude of Gibbs phenomena for $u = \pm 10\ m/s$ impact. (b) Width of the shock for $u = \pm 10\ m/s$ impact. (c) Amplitude of Gibbs phenomena for $u = \pm 100\ m/s$ impact. (d) Width of the shock for $u = \pm 100\ m/s$ impact. (e) Amplitude of Gibbs phenomena for $u = \pm 1000\ m/s$ impact. (f) Width of the shock for $u = \pm 1000\ m/s$ impact.
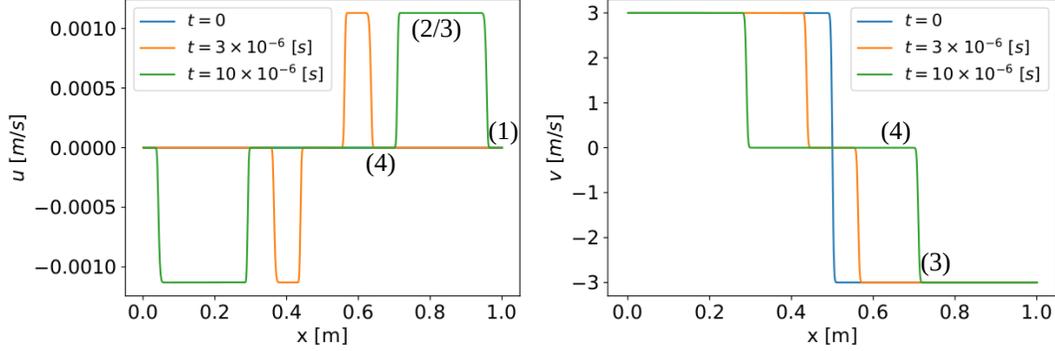
FIGURE 6.7: Wave structure for a shear Riemann problem. The jump from states (1) to (2) is an elastic normal shock. The jump from states (3) to (4) represents both shear and normal deformation.

To derive the jump condition for $\mathbf{g}$, it is helpful to consider its governing equation in index notation, in conservative form, with the curl terms omitted: $\partial g_{ij}/\partial t + \partial/\partial x_j \left(g_{ik} u_k \right) = 0$. Given the assumptions about 1D deformation in the $x - y$ plane, this ensures that $\underline{\mathbf{g}} = \mathbf{1}$ except for g_{11} and g_{21}, and yields two jump conditions:

$$g_{11}^{(1)} u_1 = g_{11}^{(2)} u_2 \tag{6.62}$$

$$g_{21}^{(1)} u_1 + v_1 = g_{21}^{(2)} u_2 + v_2 \tag{6.63}$$

The structure of $\underline{\mathbf{g}}$ guarantees that $\det \underline{\mathbf{g}} = g_{11} = \rho/\rho_0$, so the first jump condition is redundant with conservation of mass.

To solve for the states on either side of an elastic shear shock, it is helpful first to illustrate the wave structure in an elastic solid generated by a Riemann problem where the initial condition is quiescent except for a discontinuity in transverse velocity. A highly resolved example in the same copper-like stiffened gas used elsewhere in this work is provided in Figure 6.7. This initial condition generates two shocks. The first (jump 1-2) is an elastic normal shock, which travels at the longitudinal wave speed for the material ($c_{s,\perp} = \sqrt{(\gamma p_\infty + 4/3\mu)/\rho}$). Behind it is a normal/shear shock (jump 3-4), which travels more slowly, at the shear wave speed $c_{s,//} = \sqrt{4/3\mu/\rho}$. States 2 and 3 are the same in terms of internal variables, but because the shocks travel at different speeds, it is convenient to treat them separately.

Solving for the shock states must be done simultaneously by augmenting the system of equations for normal shocks, which had 5 equations and 6 unknowns. The jump relations for the shear shock add 5 additional equations, and the EOS and constitutive equation add 2 more. There are 9 new unknowns: $\rho_4, p_4, u_3, u_4, v_3, v_4, e_4, \sigma_{11}^4, \sigma_{12}^4$, which results in a total of 12 equations and 15 unknowns. To close the system, additional constraints in the velocities are used. We choose $v_4 = 0$ for simplicity,

which amounts to choosing a reference frame moving parallel to the shear shock with the speed of the shear. In addition, we use the condition that the material is at rest in state 4: $(u_1 - u_2 = u_4 - u_3)$. This system of 14 equations and 15 is closed by setting $\Delta V = v_4 - v_3$. This system can be solved numerically by guessing strength of the normal shock (1-2), and iterating until a solution consistent with the overall shear is obtained.

For the first set of tests, we consider a shear elastic precursor (jump 3-4) for the 3 materials, and the code is run in purely elastic mode, similar to the procedure used for elastic normal shocks. In these tests, we assess the effect of the coefficient in the artificial diffusivity for $\underline{\mathbf{g}}^e$, C_{ge}, while maintaining $C_\kappa = 0.01$ and $C_\beta = 5.0$, based on the previous results on normal shocks. We emphasize that the switching function employed in g^{e*} avoids adding extra dissipation to normal shocks, which allows us to tune C_{ge} and C_β separately.

In Figure 6.8 we show the effect of varying C_{ge} on the elastic (deviatoric) shear stress σ'_{21} and the artificial diffusivity for $\underline{\mathbf{g}}^e$, g^{e*}. This shows that jumps in shear stress can be effectively regularized for a variety of materials, and that the artificial diffusivity is localized to the jump in shear stress. The effect of C_{ge} on the Gibbs phenomena around the shock and the shock width is shown in Figure 6.9. Compared to normal shocks, there is a greater dependence on the material properties: the weaker materials (lower yield stress) require a larger C_{ge} to reduce the amplitude of the oscillations. This is because the normal velocity jump in a shear shock is much smaller than in a normal shock. As a result, the mechanism thinning the shock is weaker, and the shock thickness is more sensitive to the artificial diffusion added. Given that the focus of this work is on problems involving strong shocks which cause substantial plastic deformation, we use $C_{ge} = 1$ in most of the test cases that follow. For problems where the elastic deformation is more important than plastic deformation, a larger C_{ge} is suggested. For problems involving little shear deformation, a smaller value of C_{ge} may be sufficient.

If the shear is so strong that the material yields, then the resulting deformation is fluid-like, and a shock does not form, in contrast to the plastic normal shock. Consequently, we we do not simulate it in a moving reference frame. Instead, we impose a shear velocity jump ΔV on top of the elastic precursor state (state 4, when it lies just on the yield surface). Building on the previous sections we use $C_\kappa = 0.01$, $C_\beta = 5$, $C_{ge} = 1$, and in this section use plastic shear discontinuities to tune C_{gp}. In this section we consider only perfect plasticity, because it represents the greatest challenge in terms of simulation stability. Compared to a perfectly plastic model, strain hardening will tend to decrease the strength of jumps in flow variables, with the energy associated with shocks going into work-hardening.

Because the plastic shear discontinuity is a visco-plastic discontinuity, and not a shock, the nature of the regularization needed is also different. This straining field results in expansion at the discontinuity: ρ decreases and $\nabla \cdot \boldsymbol{u} > 0$. As a result, the components of $\underline{\mathbf{G}}^p$ take on a delta-function shape around the discontinuity. Without some amount of regularization, the delta function will
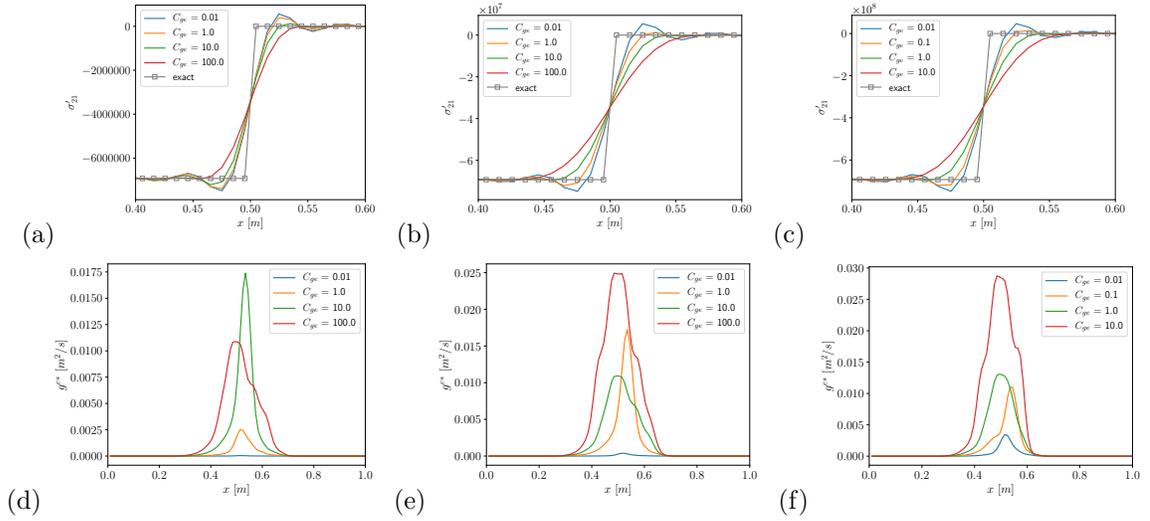
FIGURE 6.8: Regularization of elastic shear shocks as C_{ge} is varied. The sub-figures show the effect on (a) pressure for $\sigma_Y = 0.1\sigma_{cu}$, (b) pressure for $\sigma_Y = \sigma_{cu}$, (c) pressure for $\sigma_Y = 10\sigma_{cu}$, (d) artificial bulk viscosity for $\sigma_Y = 0.1\sigma_{cu}$, (e) artificial bulk viscosity for $\sigma_Y = \sigma_{cu}$, and (f) artificial bulk viscosity for $\sigma_Y = 10\sigma_{cu}$.
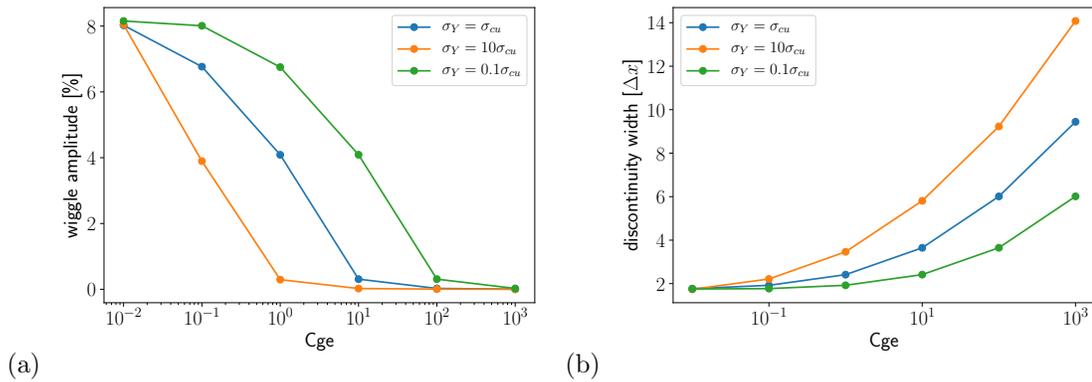
FIGURE 6.9: Effect of C_{ge} in elastic shear precursor shocks in 3 different materials. (a) Amplitude of Gibbs phenomena around the shock. (b) Width of the shock.
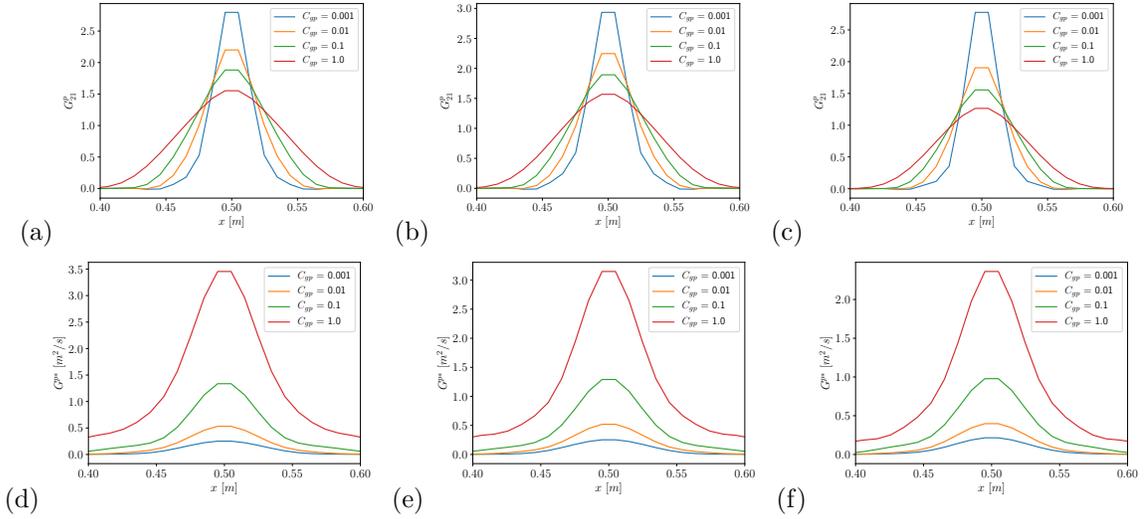
FIGURE 6.10: Regularization of plastic shear discontinuity as C_{gp} is varied for an overall shear of $v = \pm 100 \ m/s$. The sub-figures show the effect on (a) G_{21}^p for $\sigma_Y = 0.1\sigma_{cu}$ (b) G_{21}^p for $\sigma_Y = \sigma_{cu}$ (c) G_{21}^p for $\sigma_Y = 10\sigma_{cu}$ (d) artificial $\underline{\mathbf{G}}^p$ diffusivity for $\sigma_Y = 0.1\sigma_{cu}$ (e) artificial $\underline{\mathbf{G}}^p$ diffusivity for $\sigma_Y = \sigma_{cu}$ (f) artificial $\underline{\mathbf{G}}^p$ diffusivity for $\sigma_Y = 10\sigma_{cu}$.

lead to Gibbs phenomena, which grow unbounded. In this section, only values of C_{gp} which keep the oscillations bounded are shown. This different shape of the discontinuity also means that the interface width and wiggle amplitude must be quantified differently than in the preceding sections. For wiggle amplitude, we use the height of the Gibbs phenomena to the undeformed value divided by the height of the delta function. For interface width, we use the 95% thickness, where the oscillations about the undeformed value have decayed to less than 5% of the height of the jump. Finally, since this initial condition does not reach a steady state, we also assess the discontinuity width and wiggle amplitude at an earlier time. The simulation is run until a wave traveling at the speed of sound can traverse the domain 5 times, so that initial transients vanish, but the spreading due to the expansion is kept small.

Figure 6.10 shows the influence of C_{gp} on the regularization of the discontinuity in $\underline{\mathbf{G}}^p$, focusing on the G_{21}^p component, which most directly quantifies shear. As C_{gp} is increased, the delta function is spread over more grid points. There are small differences depending on the material strength, mainly that with the larger yield, a larger coefficient is needed to achieve the same degree of spreading. With even the smallest stable value of C_{gp}, the discontinuity is regularized quite well, without visible Gibbs phenomena. This is confirmed in Figure 6.11, which shows the effect on Gibbs oscillation amplitude and delta function width quantitatively. These results suggest that only a very small value of C_{gp} – in the range of 10^{-3}-10^{-2} – is needed for regularization. Note that the wiggle amplitudes for the weaker shear can be deceptive, as the height of the delta function is small, and the metric is sensitive to this denominator.
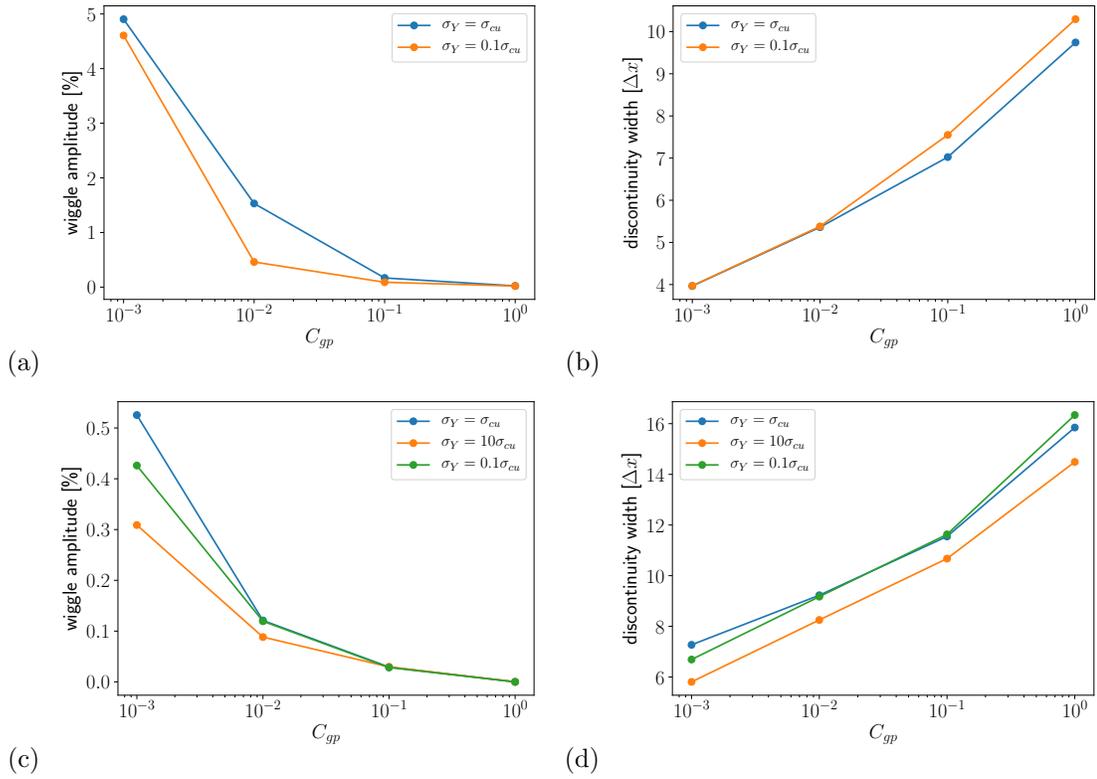
FIGURE 6.11: Effect of C_{gp} in elastic-plastic normal shocks in 3 different materials and 2 different shear magnitudes. (a) Amplitude of Gibbs phenomena for $v = \pm 10 \ m/s$ shear. (b) Width of the delta function discontinuity for $v = \pm 10 \ m/s$ shear. (c) Amplitude of Gibbs phenomena for $v = \pm 100 \ m/s$ shear. (d) Width of the delta function discontinuity for $v = \pm 100 \ m/s$ shear.
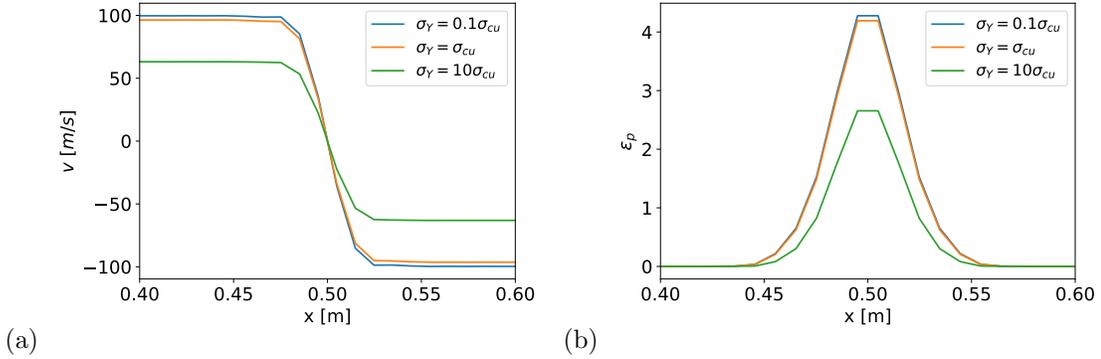
FIGURE 6.12: (a) Velocity and (b) equivalent plastic strain profiles for shear magnitude $v \pm 100 \ [m/s]$ with $C_{gp} = 0.01$.

Since the plastic shear test was carried out with perfect plasticity, the results do not couple back to the velocity field, and the velocity field is responsible for the thickness of the delta function in $\underline{\mathbf{G}}^p$ with very small C_{gp}. This is shown in Figure 6.12, which uses $C_{gp} = 0.01$. This shows that the transverse velocity v discontinuity is effectively regularized with artificial shear viscosity and artificial diffusivity of $\underline{\mathbf{g}}^e$. Because the discontinuity is spread over approximately 5 grid points with no visible Gibbs phenomena, this explains why the profiles in $\underline{\mathbf{G}}^p$ are so smooth. In simulations with strain hardening, the regularization of $\underline{\mathbf{G}}^p$ with provide additional smoothing through coupling back to the shear modulus and yield stress. In additional, we show in Figure 6.12(b) that the Euler-Almansi plastic strain (equivalent plastic strain) is effectively regularized.

Because the plastic shear test does not reach a steady state, and thus is prone to becoming thicker than other discontinuities in this 1-D setting, we finish the evaluation of g^{p*} to regularize plastic shear by examining how rapidly the discontinuity widens. The discontinuity width is plotted against time for the case of $v = \pm 100 \ [m/s]$ with nominal copper properties in figure 6.13. A typical diffusion process would cause the discontinuity to widen as $\Delta \sim t^{1/2}$. However, the observed power law is closer to $\Delta \sim t^{1/4}$, meaning the LAD approach spreads the discontinuity more slowly than a constant diffusivity.

6.3.3   1-D impact tests

To demonstrate the code capability in transient problems with shocks, we consider one-dimensional impact problems of two semi-infinite slabs. The test problem used here is the same as that used in section 4(a) by Favrie and Gavrilyuk [59]. We also extend this problem to demonstrate shock propagation with more complex material behavior by adding strain hardening.

The material properties used are $\gamma = 2.84$, $p_\infty = 0.6 \ GPa$, $\mu = 77 \ GPa$, and $\sigma_Y = 2.49 \ GPa$. The initial pressure is 0.1 MPa. For perfect plasticity, σ_Y is constant, but for the strain hardening
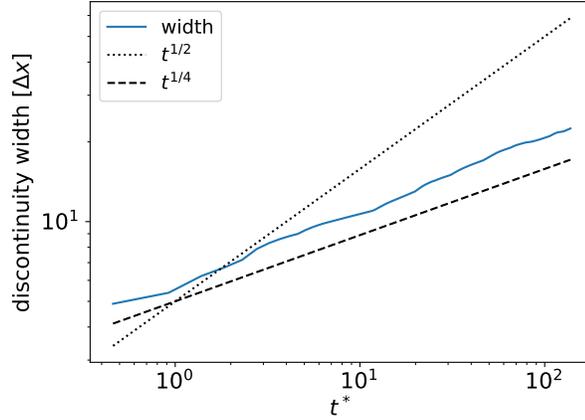
FIGURE 6.13: Discontinuity width over time for the case $v \pm 100 \; [m/s]$ with nominal copper properties, with $C_{gp} = 0.01$. Time is non-dimensionalized by the domain length and the speed of sound.

we consider a simplified version of the Steinberg-Cochran-Guinan model [186]: $\sigma_Y = \sigma_{Y,0}(1+\beta\epsilon_p)^n$. We use the parameters $\beta = 36$ and $n = 0.45$, which are the values suggested for copper. We consider symmetric impacts with $u = \pm 50$, 150, and 300 m/s, and across all tests in this section $C_\beta = 5$ is used. Figure 6.14 shows profiles of u, ρ, and β^* after $t = 0.01 \; sec$. For the impact velocity $u = \pm 50 m/s$, the material does not yield, and only an elastic wave is generated. For the two stronger impacts, both an elastic precursor and an elastic-plastic shock are created. The elastic-plastic shock travels much more slowly. This is expected because the shear modulus no longer contributes to the sound speed [132].

The main difference with the addition of strain hardening is that the elastic-plastic shock involves less compression (smaller density jump) and travels significantly faster. The faster speed can be explained by rearranging the jump condition for mass in a stationary frame of reference:

$$U_s = u_1 - \Delta u \frac{1}{\frac{\rho_1}{\rho_2} - 1} \tag{6.64}$$

where U_s is the shock speed, and $\Delta u = u_2 - u_1$. Because the elastic precursor is the same regardless of strain hardening, u_1 and Δu are unchanged, and the shock speed only depends on ρ_1/ρ_2. Since strain hardening reduces the amount of compression, this explains the increase in speed. Intuitively, this can be explained because with strain hardening, the elastic-plastic shock which follows the precursor involves additional elastic deformation relative to perfect plasticity, and the speed of elastic waves exceeds plastic waves [132].

The speeds of the elastic and plastic waves with and without strain hardening are shown in an X-T diagram in Figure 6.15. The shock locations are extracted from simulation data by detecting local extrema in du/dx, which results in some initial inaccuracy near $t = 0$. All the shock trajectories
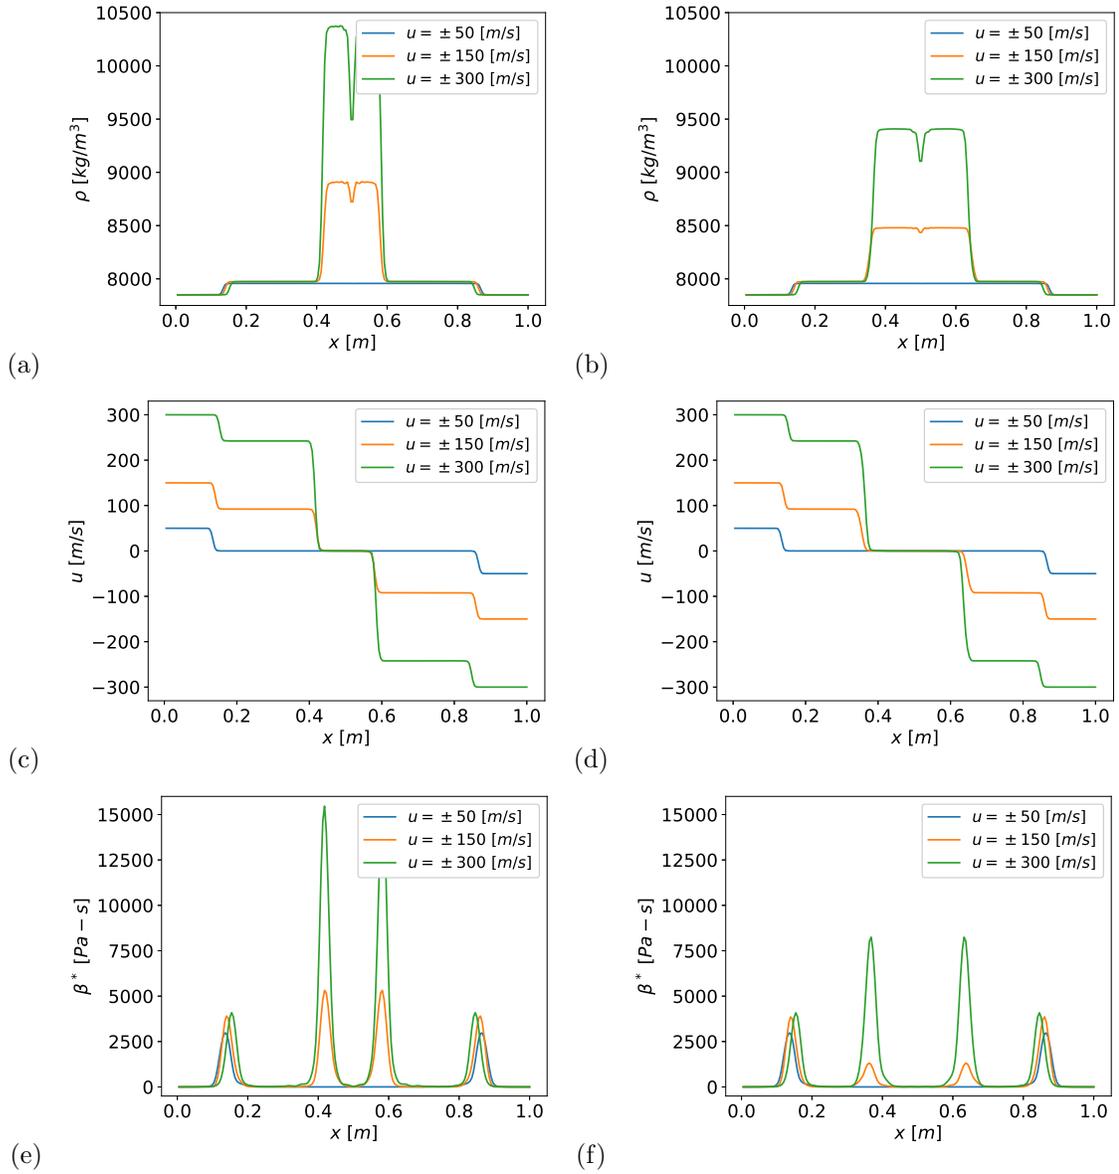
FIGURE 6.14: One-dimensional impact problem. (a) Density with three different impact strengths and perfect plasticity. (b) Density with three different impact strengths and simple strain hardening. (c) Axial velocity with three different impact strengths and perfect plasticity. (d) Axial velocity with three different impact strengths and simple strain hardening. (e) Artificial bulk viscosity with three different impact strengths and perfect plasticity. (f) Artificial bulk viscosity with three different impact strengths and simple strain hardening.

(a)   (b)
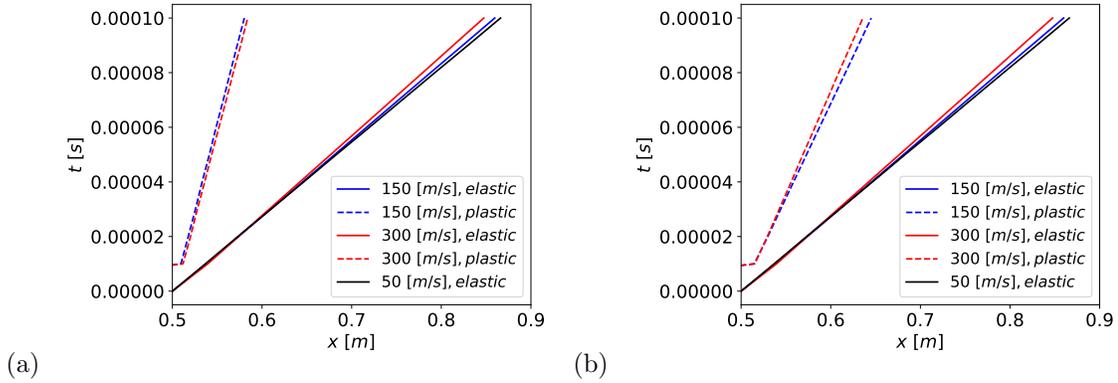
FIGURE 6.15: *X-T* diagram of elastic and plastic shocks in an impact problem with (a) perfect plasticity and (b) strain hardening.

Shock type	Plasticity model	Δu_0 [m/s]	Observed speed	Predicted speed	Relative error
elastic	perfect	50	3658.5	3659.1	-0.016%
elastic	strain hardening	50	3658.5	3659.1	-0.016%
elastic precursor	perfect	150	3576.3	3569.0	0.205%
elastic precursor	strain hardening	150	3576.8	3569.0	0.220%
elastic precursor	perfect	300	3413.9	3419.0	-0.147%
elastic precursor	strain hardening	300	3414.0	3419.0	-0.146%
plastic	perfect	150	785.3	783.8	0.187%
plastic	strain hardening	150	1443.8	1434.9	0.620%
plastic	perfect	300	800.1	806.4	-0.779%
plastic	strain hardening	300	1334.9	1345.1	-0.756%

TABLE 6.1: Shock speeds in 1d impact problem with and without strain hardening. Observed shock speed is estimated from a least squares fit to the last 5 simulation visualization files. Predictions are from a numerical jump-conditions solver implemented in MATLAB. Relative error is relative to the predicted shock speed.

are linear, indicating constant speeds with or without strain hardening. The speed of shocks with or without strain hardening can also be predicted from the jump conditions described earlier. To solve for a shock speed with strain hardening, the solution procedure must be modified to account for the dependence of the yield stress on the plastic deformation via $\underline{\mathbf{G}}^p$. The observed shock speeds are compared to the predicted shock speeds in Table 6.1. The observed behavior corresponds well with the predictions. In particular, the independence of the elastic precursor speed on strain hardening, the increase in plastic shock speed with strain hardening, and the decrease in elastic precursor speed as overall impact velocity is increased, are all captured quantitatively.

A grid refinement exercise for the strongest impact is shown in Figure 6.16, which shows that both cases converge and result in more sharply refined shocks as the grid is refined. As expected, the artificial bulk viscosity required to regularize the shock becomes smaller with grid refinement.
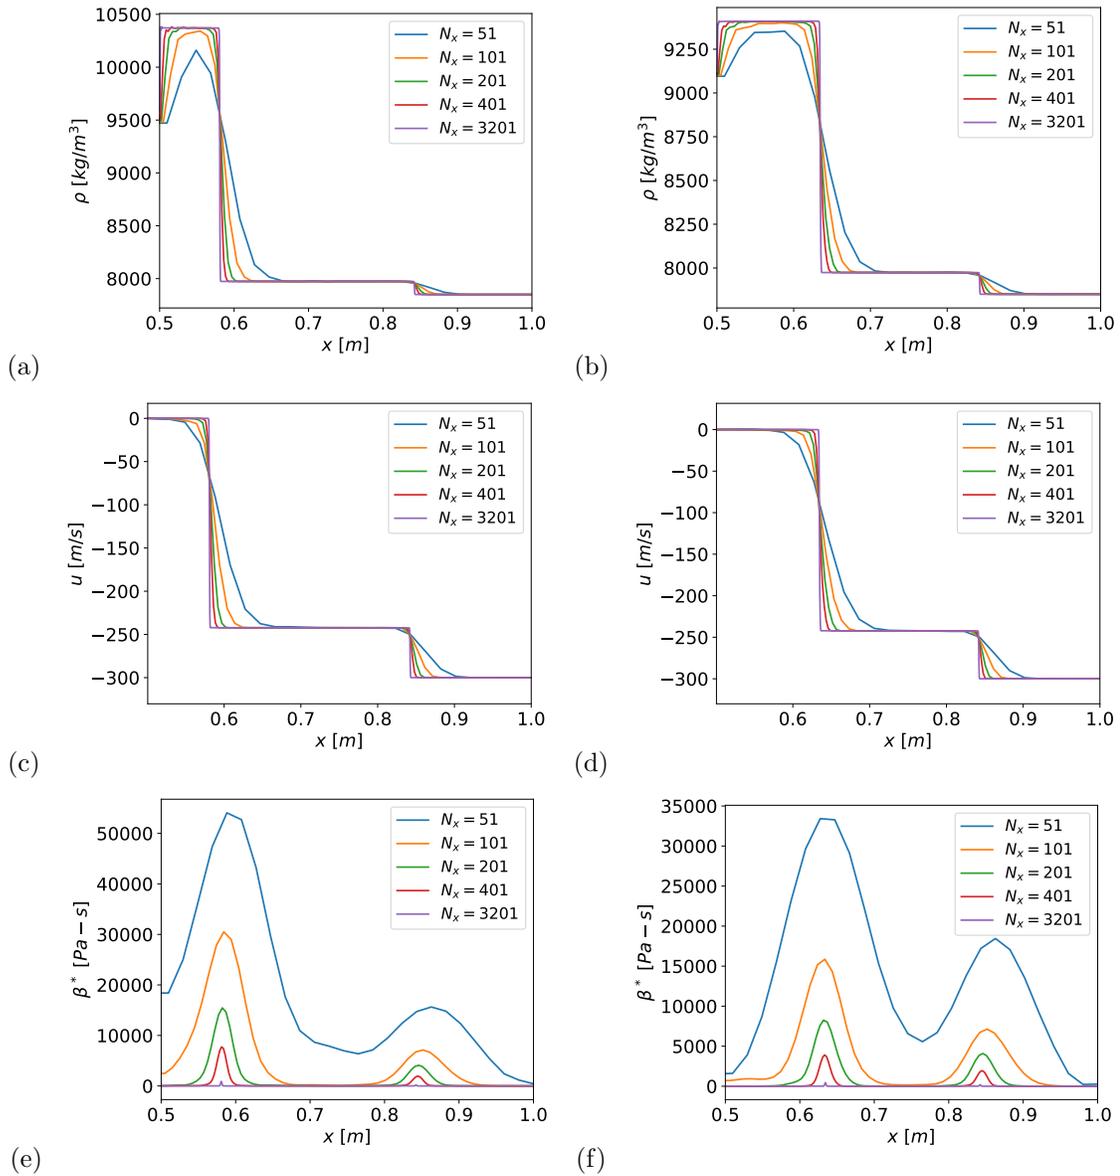
FIGURE 6.16: Grid refinement study in a one-dimensional impact problem (a) Density with perfect plasticity (b) Density with strain hardening (c) Axial velocity with perfect plasticity (d) Axial velocity with strain hardening (e) Artificial bulk viscosity with perfect plasticity (f) Artificial bulk viscosity with strain hardening.

6.3.4 Multiscale resolution capability

In this section we introduce a new test problem, inspired by the Shu-Osher problem [178], the purpose of which is to assess the ability of the numerical method to capture fine-scale features interacting with elastic and elastic-plastic shocks without adding excessive dissipation. For simplicity of definition, we modify the preceding 1D impact test case from Favrie and Gavrilyuk [59]. Using the same material and initial condition, we superimpose a density perturbation onto the initial state of the material. We also extend the domain to $x = [0, 10]m$ to give better separation between the elastic precursor and plastic shocks. In the traditional Shu-Osher problem, the initial pressure is uniform, and the density perturbation also corresponds to variations in internal energy. In the solid, instead of constant pressure, constant normal stress (σ_{11}) is required for static equilibrium at the initial condition. This is achieved by assuming the density perturbation corresponds to uniaxial elastic deformation: $g_{11}^e = \rho/\rho_0$. From this, the deviatoric (elastic) stress can be calculated, and the pressure required for static equilibrium is solved for. In equation form, the initial condition is

$$u(x, t_0) = 300\big(1 - 2\mathcal{H}(x - 5)\big) \, [m/s] \tag{6.65}$$

$$\rho(x, t_0) = 7850 \left(1 + 0.001 \cos \frac{200\pi}{10}(x - 5) \right) \, \left[\frac{kg}{m^3} \right] \tag{6.66}$$

$$\underline{\mathbf{g}}^e(x, t_0) = \begin{pmatrix} \rho(x, t_0)/\rho_0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \tag{6.67}$$

$$\underline{\mathbf{g}}^p(x, t_0) = \mathbf{1} \tag{6.68}$$

$$p(x, t_0) = p_0 + \sigma'_{xx}\big(\underline{\mathbf{g}}^e(x, t_0)\big) \tag{6.69}$$

where $\mathcal{H}(x)$ is the Heaviside function, and $p_0 = 10^5 \, Pa$. The density perturbation specified here is small enough so that the resulting elastic stress does not cause the material to yield.

To visualize the wave structure in this problem, "carpet plots" of density and velocity for the right half of the domain are shown in Figure 6.17. In these plots, profiles are offset vertically in proportion to the simulation time. The elastic precursor and plastic shock are features from the original problem which remain. The density perturbation is preserved through the elastic precursor, and after the passage of the plastic shock, it is significantly amplified. The other notable feature is asymmetrical density perturbations in the region between the shocks. These correspond to local plastic deformation, which occurs only in the crests of the density perturbation. This occurs because of the elastic stress which existed in the material before the elastic precursor shock. These asymmetric perturbations grow until they are swallowed up by the plastic shock, after which the asymmetry is "forgotten" and a larger amplitude wave is left. This series of density perturbation corresponds to a pattern of small plateaus and compressions in velocity, shown in Figure 6.17(b).

(a)   (b)
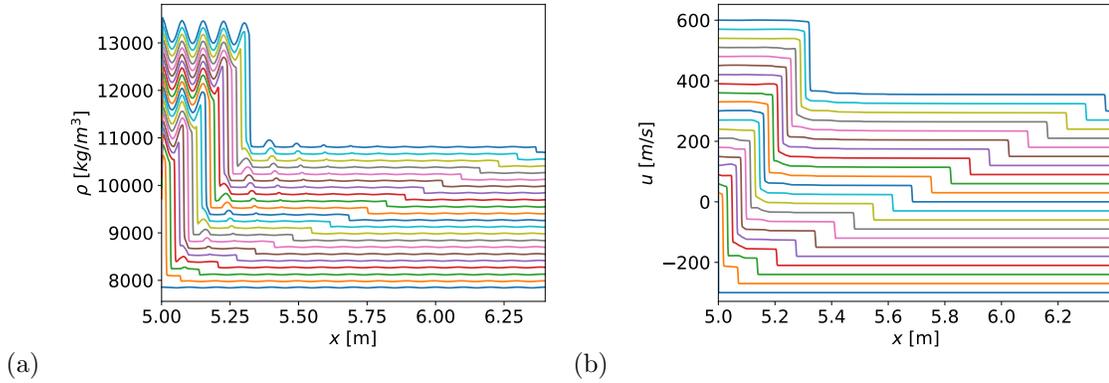
FIGURE 6.17: Time evolution of a Shu-Osher style impact problem in (a) density and (b) velocity. Profiles are offset vertically proportional to the simulation time to aid in the visualization of waves. The reference solution using $N_x = 10,001$ is plotted.

The effect of grid resolution while using the recommended LAD coefficients ($C_\beta = 5$, $C_\kappa = 0.01$) on certain solution features is shown in Figure 6.18. For the coarsest case ($N_x = 501$) there are not sufficient grid points to represent the wave, but for higher resolutions, the amplitude of the perturbations is consistently predicted through the different shocks. At higher resolution, the shocks are thinned and Gibbs phenomena are reduced, particularly those which arise around the edges of the asymmetric density perturbations.

The effect of C_β is shown in Figure 6.19. Around the elastic precursor shock, shown in Figure 6.19(a), we see that if C_β is too low and the Gibbs oscillations are too strong, there are amplitude errors after the shock, but that these are very small for $C_\beta \geq 1$. Smaller values of C_β also lead to amplitude errors ahead of the plastic shock, and to phase errors behind it. Increasing C_β too far does not appear to have any negative consequences beyond excessive spreading of the shock.

The effect of C_κ is shown in Figure 6.20. It has been previously demonstrated by Cook [39] that excessive artificial heat conduction can attenuate temperature oscillations after the shock. This is the case here, but for values of $C_\kappa \leq 0.01$, the difference is negligible. The attenuation is present around both elastic precursor and plastic shocks, but it is more severe around the plastic shock.

Based on this test problem, the values of $C_\beta = 5$ and $C_\kappa = 0.01$ are still recommended for elastic-plastic deformation in metals. Further, this test has shown that despite the lower formal order of temporal accuracy associated with the implicit plastic relaxation, this has not been detrimental to the resolution characteristics of the high order compact schemes, which give minimal phase and amplitude error in this type of problem.
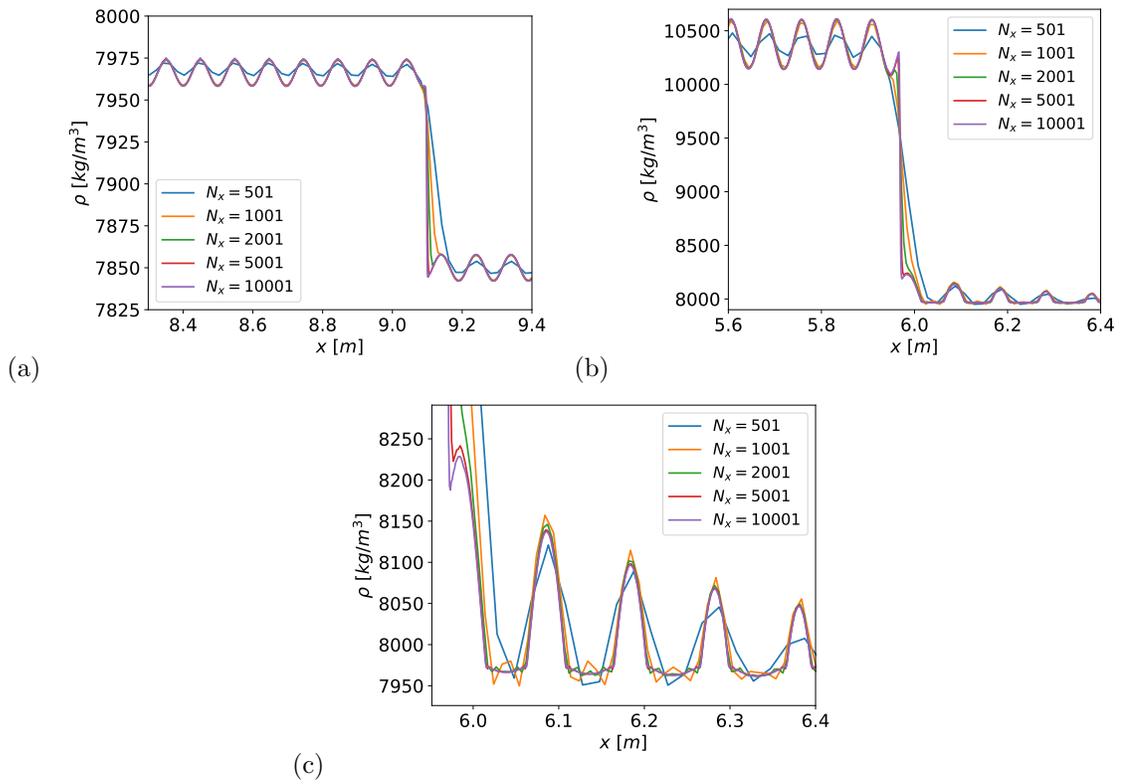
FIGURE 6.18: Effect of grid resolution on density in the Shu-Osher style impact problem. Panels highlight (a) the elastic precursor, (b) the plastic shock, and (c) the density perturbations in front of the plastic shock.
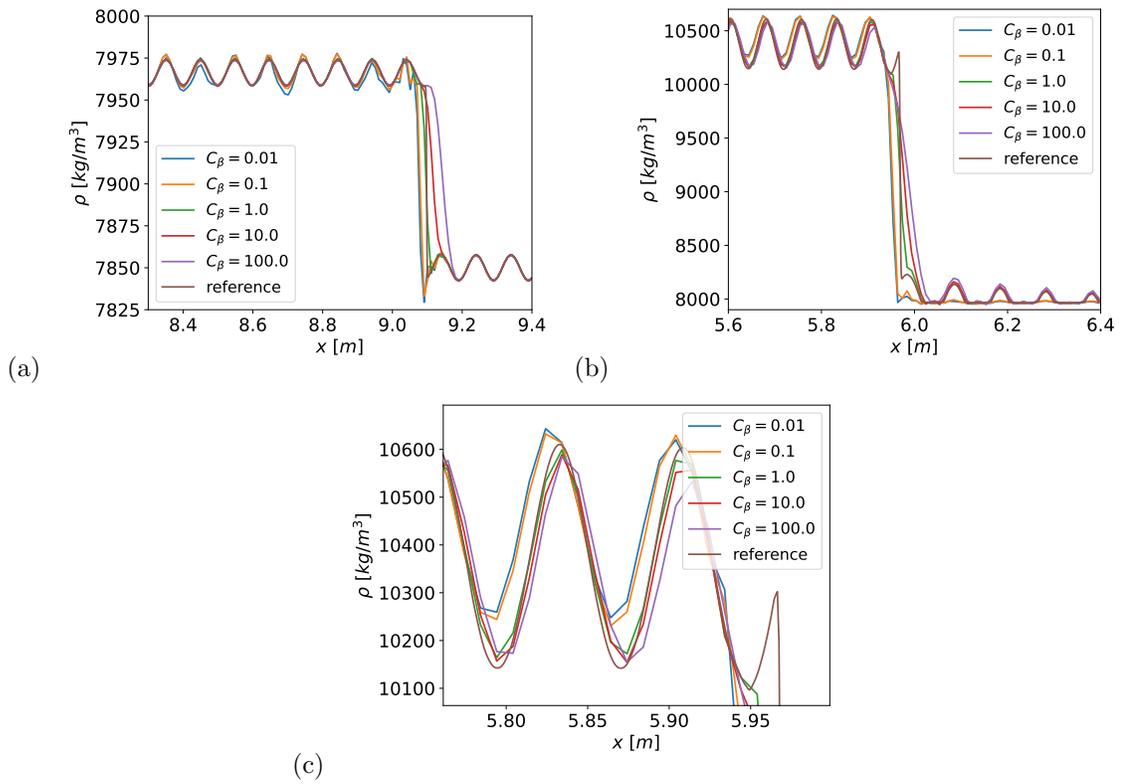
FIGURE 6.19: Effect of C_β on density in the Shu-Osher style impact problem. Panels highlight (a) the elastic precursor, (b) the plastic shock, and (c) the waveform after the passage of the plastic shock.
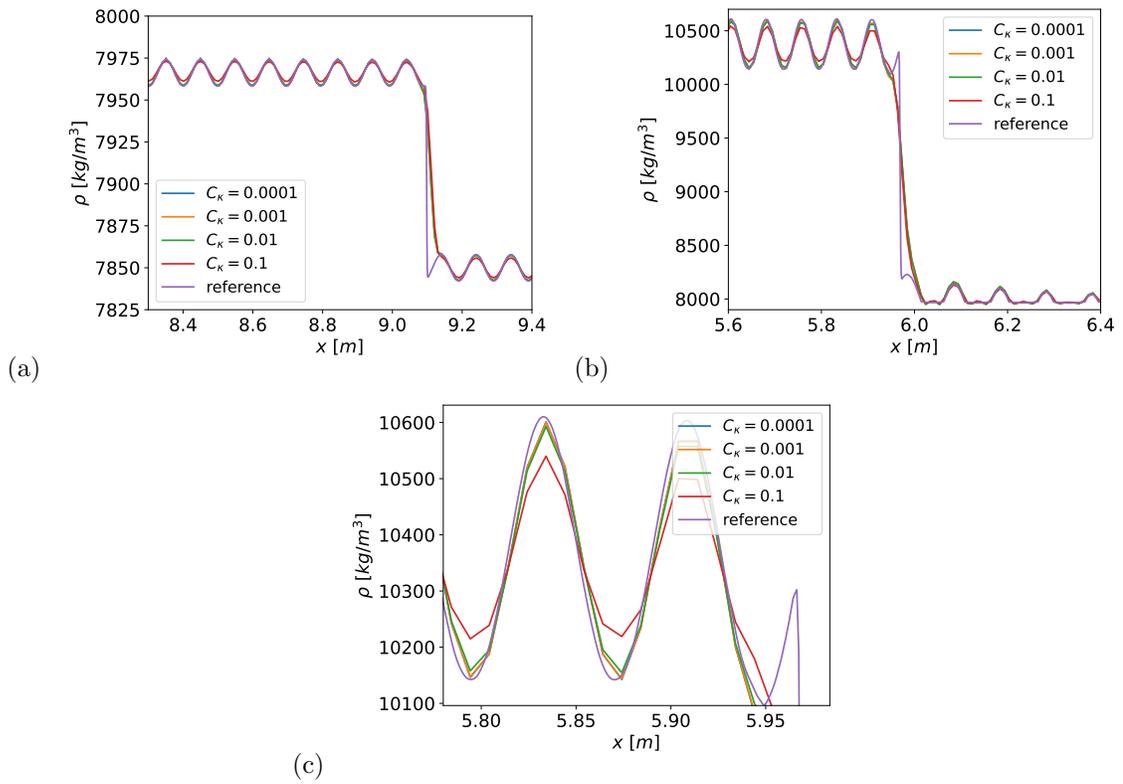
FIGURE 6.20: Effect of C_κ on density in the Shu-Osher style impact problem. Panels highlight (a) the elastic precursor, (b) the plastic shock, and (c) the waveform after the passage of the plastic shock.

6.4 Taylor impact

The Taylor impact problem was originally proposed as a means to estimate the dynamic yield stress of materials by Taylor [192]. It is an important problem for assessing the accuracy of strain hardening models and simulation codes because simulation results can be directly compared with experimental measurements of the resulting shape. In this section, we simulate a Taylor impact problem and compare to experimental data to demonstrate the capability of the method for handling problems with strong, elastic-plastic shocks in which strain hardening plays a significant role. Further, we demonstrate capability with different strain hardening models, including rate-dependent plasticity.

The impact problem we consider is case "Cu-1" from Banerjee [14], which corresponds to an experiment from Wilkins and Guinan [208]. In this test, a cylinder of OFHC copper which is 23.47 mm long and 7.62 mm in diameter impacts a rigid wall at 210 m/s, at a temperature of 298 K. To simulate this problem with the framework used here requires a few approximations. First, the boundary condition at the impact plane is implemented here as a symmetry boundary condition. In physical terms, this would correspond to two cylinders which collide and fuse together. The simulation is initialized with the copper cylinder in contact with the symmetry boundary condition. All together, this amounts to neglecting the physics of the lubrication flow beneath the cylinder, as well as the friction between cylinder and wall. A second approximation is used in treating the ambient fluid. Because copper and air have very different compressibilities, this can lead to difficulties in the mixture equation of state, which employs an iterative solver to achieve pressure and temperature equilibrium. Pressure-temperature equilibrium can also lead to problems with wave transmission through material interfaces due to a non-monotonic mixture speed of sound [175]. To circumvent these issues, the air is replaced in this problem with a stiffened gas which has the density of air but a large stiffening pressure. The properties of the ambient fluid are $\rho_0 = 1.225 \ kg/m^3$, $\gamma = 1.4$, $p_\infty = 0.195 \ GPa$. Because the air is mostly a non-participant in the deformation of the cylinder, these approximations still result in meaningful results, as shown below.

The copper is represented using a Mie-Grüneisen equation of state, with the coefficients used in Banerjee [14]. Two widely-used plasticity models are considered: Steinberg-Cochran-Guinan [186] and Johnson-Cook [102]. Based on the preceding sections, and the knowledge that this problem does not involve significant shear deformation, the LAD coefficients that we use for this problem are $C_\beta = 5.0$, $C_\kappa = 0.01$, $C_\mu = 0.002$, $C_{ge} = 0.1$, and $C_{gp} = 0.001$.

Finally, we note that in this section, we use the same numerical method as before, implemented into the Miranda code, developed at Lawrence Livermore National Laboratory (LLNL). This is done because of the availability of cylindrical coordinates, and the flexibility with a variety of equation of state and strain hardening models. Discretizing the evolution equations for $\underline{\mathbf{g}}^e$ and $\underline{\mathbf{G}}^p$ involves computing differential operators of high-rank tensors, including the gradients of a rank-2 tensor and divergence of a rank-3 tensor. The details of these operators in Cartesian, cylindrical, and spherical coordinates are provided in Adler et al. [2]. In this section, we use cylindrical coordinates
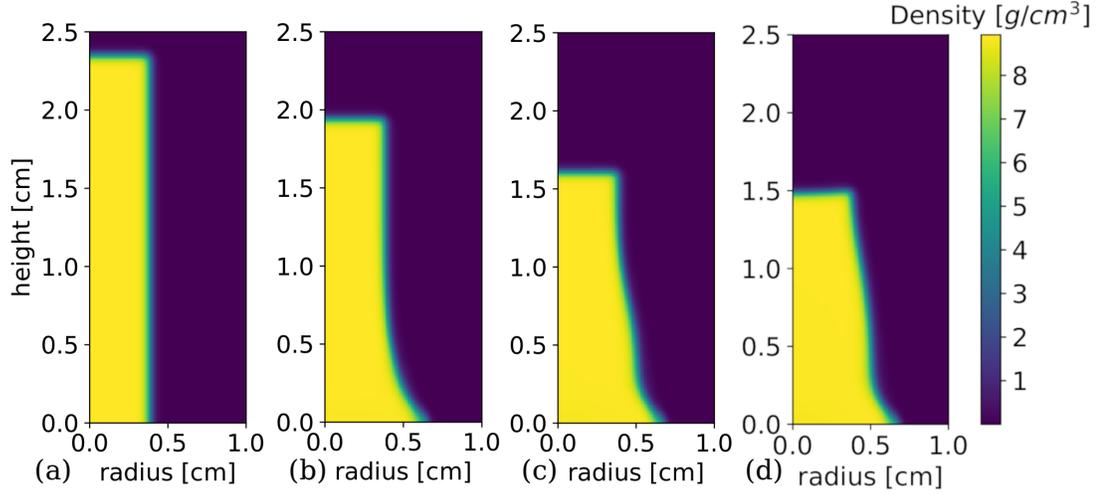
FIGURE 6.21: Snapshots of mixture density in a Taylor impact problem with the Steinberg-Cochran-Guinan strength model at times (a) $t = 0$, (b) $t = 20 \times 10^{-6}$ seconds, (c) $t = 40 \times 10^{-6}$ seconds, and (d) $t = 60 \times 10^{-6}$ seconds. The domain is truncated for plotting purposes.

in axisymmetric mode to minimize computational cost. The Miranda version of the code also uses the "single-g" formulation, in which a single $\underline{\mathbf{g}}^e$ and $\underline{\mathbf{G}}^p$ are solved for the multiphase mixture.

A simulation domain of $L_r = 3cm$, $L_z = 3cm$ is used, with uniform grid defined by $N_r = 360$, $N_z = 360$. Finer and coarser resolution cases (180 and 720) are considered to assess the effect of the grid. The material interface is initialized over $\mathcal{O}(8)$ grid points. The simulation is run for 100×10^{-6} seconds, which is sufficient time for the cylinder to come to rest in its final shape, although it continues to vibrate in place as elastic waves bounce around inside it.

A time series of the simulation using the Steinberg-Cochran-Guinan (SCG) strength model with $N_r = 360$ is shown in Figure 6.21. The last snapshot shown is for $t = 60 \times 10^{-6}$ seconds, but subsequent snapshots are nearly identical. The density inside the copper slug is nearly constant, while the shape changes substantially due to plastic deformation. The material interface does not thicken significantly during the simulation.

As the grid is refined, the material interface and plastic deformation are both more sharply resolved. This is illustrated in Figure 6.22, which shows contours of equivalent plastic strain and the 50% volume fraction contour. At lower resolution, it is noticeable that there is a small amount of plastic strain which accumulates in regions of small copper volume fraction, but this is damped according to the method described in Section 6.2.9. As the grid is refined, additional plastic strain is resolved, and that strain is more concentrated at the origin of the cylinder. These contours agree qualitatively with previously reported results in Barton and Romenski [19], which show a local minimum in plastic entropy around the same distance from the wall, in addition to a sharp gradient in plastic strain at the bottom of the cylinder.
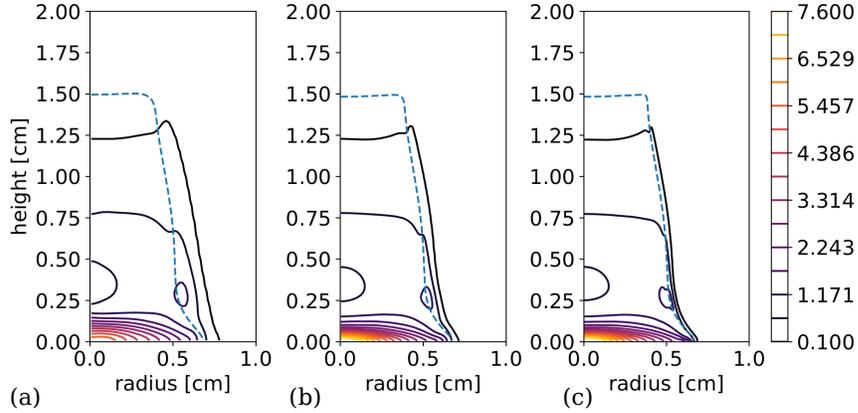
FIGURE 6.22: Contours of equivalent plastic strain under grid refinement with the Steinberg-Cochran-Guinan plasticity model. (a) $N_r = 180$, (b) $N_r = 360$, (c) $N_r = 720$. The 50% volume fraction contour is plotted with a dashed blue line. The domain is truncated for plotting purposes.

Finally, we demonstrate the capability of the method to use rate-dependent plasticity models by comparing the final shape from simulations using the SCG model and the Johnson-Cook (JC) model in Figure 6.23. In both cases, the overall cylinder shape changes very little with grid refinement, except for the upper right corner becoming sharper, indicating the results are very nearly converged. The differences in shape between the two models are expected, and agree qualitatively with the comparison done by Banerjee [14]. he Johnson-Cook model predicts the spreading of the base more accurately, while slightly under-predicting the bulging away from the wall. In contrast, the SCG model under-predicts spreading of the bottom and over-predicts bulging away from the wall. Neither model converges perfectly to the experimental results, which is expected given their empirical nature and the approximation made in representing the boundary condition as a symmetry boundary. Finally, we note that the agreement of the Johnson-Cook prediction with experiment is not as good as that shown in Banerjee [14]. This may be due to differences in the choice of constitutive equation (stress-strain relation) and model for the shear modulus, which vary widely even across recent studies [20, 48, 76, 157].

6.5   Solid-solid Richtmyer-Meshkov instability

In this section we demonstrate the capability of the method on problems involving significant shear deformation and distortion of material interfaces which would be challenging for Lagrangian approaches. We simulate a 2D Richtmyer-Meshkov instability (RMI) between two solids, with and without strain hardening. In this section, the solids are idealized versions of copper and aluminum, represented as stiffened gases, and a simple strain hardening model is used. This is an idealized
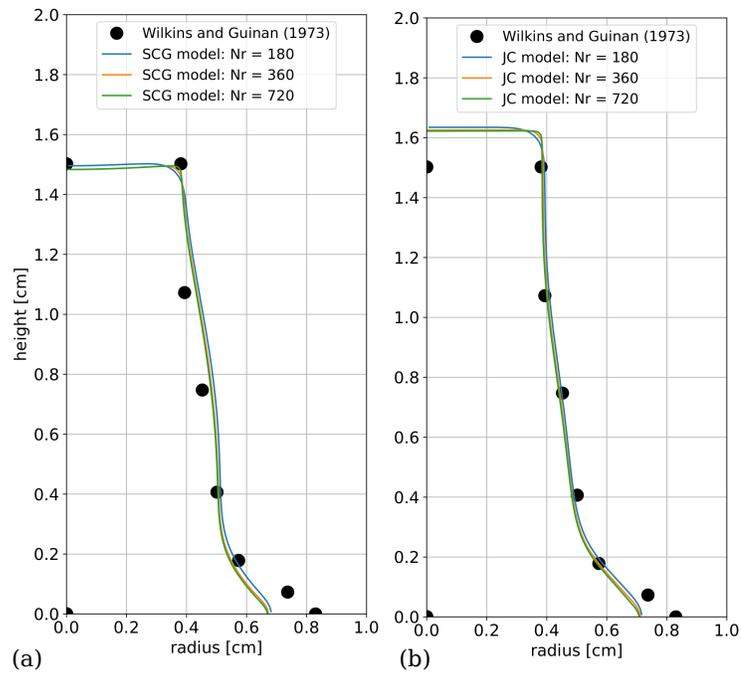
FIGURE 6.23: Final shapes of the copper cylinder in a Taylor impact problem, as visualized by the 50% volume fraction contour, for several grid resolutions, using (a) Steinberg-Cochran-Guinan (SCG) plasticity model and (b) Johnson-Cook (JC) plasticity model (JC). Digitized experimental data from Wilkins and Guinan [208] is included in both panels.

version of the instability that might develop during impact welding.

In this problem, a shock wave travels through the copper material toward a sinusoidally perturbed interface with aluminum. A version of this problem with perfect plasticity was studied by Lopez Ortega [120], who used a level-set method combined with the modified ghost-fluid method to set boundary conditions at material interfaces. The perfect plasticity version of this problem was also studied by Subramaniam et al. [188] and Adler and Lele [5], and the results presented here use an extension of the same code.

The problem domain spans $(-2 \leq x^* \leq 4; 0 \leq y^* \leq 1)$, where the dimensions have been normalized by the domain width, L_y. A symmetry boundary is applied at $x^* = 4$, representing a perfectly reflecting slip wall, and a sponge is applied over $(-2 \leq x^* \leq -1.5)$, approximating a nonreflecting free boundary. The boundary conditions in the y direction are periodic. The domain is discretized with a uniform Cartesian grid of size $N_x = 768$, $N_y = 128$. For the grid refinement and coarsening cases, two finer and two coarser grids are considered, with $N_x = [192, 384, 1536, 3072]$ and $N_y = [32, 64, 256, 512]$. The material properties for the copper medium are $\gamma_1 = 2.0$, $\rho_1 = 8930 \ kg/m^3$, $p_{\infty 1} = 62.617 \ GPa$, $\mu_1 = 47.7 \ GPa$, and $\sigma_{Y1} = 0.12 \ GPa$. The material properties for the aluminum medium are $\gamma_2 = 2.088$, $\rho_2 = 2712 \ kg/m^3$, $p_{\infty 2} = 32.654 \ GPa$, $\mu_2 = 27.6 \ GPa$, and $\sigma_{Y2} = 0.297 \ GPa$. The yield stresses, densities, and shear moduli are taken from Steinberg et al. [186], and the stiffening pressures are set to keep the initial speed of sound consistent with Lopez Ortega [120], Appendix C. For the cases with strain hardening, a simplified version of the SCG model is used, which includes the effect of strain hardening but omits all other effects involving temperature and pressure dependence: $\sigma_Y = \sigma_{Y,0} \left(1 + \beta(\epsilon_p)\right)^n$. The strain hardening properties for copper are $\beta_1 = 36$ and $n_1 = 0.45$, and the properties for aluminum are $\beta_2 = 125$ and $n_2 = 0.1$, following to Steinberg et al. [186]. The initial conditions for the velocity, pressure, volume fraction, and density are

$$u = u^{(2)} f_s + u^{(1)} \left(1 - f_s\right), \quad v = 0, \quad p = p^{(2)} f_s + p^{(1)} \left(1 - f_s\right),$$

$$\phi_1 = \phi_1^\epsilon + \left(1 - 2\phi_1^\epsilon\right) f_\phi, \quad \phi_2 = 1 - \phi_1, \quad \rho = \left(\phi_1 \rho_1 + \phi_2 \rho_2\right) \left[\rho^{(2)}/\rho^{(1)} f_s + \left(1 - f_s\right)\right],$$

$$(6.70)$$

respectively, in which the volume fraction function, f_ϕ, and the shock function, f_s, are given by

$$f_\phi = \frac{1}{2} \left(1 - \text{erf} \left\{ \frac{x - [2 + 0.4/(4\pi y) \sin(4\pi y)]}{3\Delta x} \right\}\right) \quad \text{and} \quad f_s = \frac{1}{2} \left[1 - \text{erf} \left(\frac{x - 1}{2\Delta x}\right)\right], \quad (6.71)$$

respectively, with jump conditions across the shock for velocity $\left(u^{(1)} = 0; u^{(2)} = 2605\right)$, density $\left(\rho^{(1)} = 8930, \ \rho^{(2)} = 13670\right)$, and pressure $\left(p^{(1)} = 0.1 \ MPa; p^{(2)} = 100 \ GPa\right)$.

The kinematic tensors are initialized in a pre-strained state consistent with elastic uni-axial compression associated with shock initialization, assuming no plastic deformation has yet occurred,

with

$$\underline{\mathbf{g}}^e = \begin{pmatrix} \left[\rho^{(2)} f_s + \rho^{(1)} \left(1 - f_s\right)\right] / \rho_1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \qquad \text{and} \qquad \underline{\mathbf{G}}^p = \mathbf{1}. \tag{6.72}$$

The LAD coefficients used in this problem are $C_\mu = 0.002$, $C_\beta = 5$, $C_\kappa = 0.01$, $C_D = 0.003$, $C_Y = 100$, and $C_{ge} = 1.0$. With strain hardening, $C_{gp} = 0.01$ is used, but with perfect plasticity, $C_{gp} = 0.1$ is used to stabilize the continued interface roll-up. However, with perfect plasticity, this additional diffusion of $\underline{\mathbf{G}}^p$ does not couple back to the solution.

Snapshots of the instability development with and without strain hardening are presented in Figure 6.24. The early development after the initial shock passage ($t^* \approx 1.25$) is very similar, but clear differences have emerged by the time of re-shock ($t^* \approx 1.625$), though the shock location remains consistent regardless of the plasticity model. After re-shock, the spikes and bubbles grow more rapidly, and roll-up begins around the tip of the spike. Strain hardening arrests all of these processes as larger elastic stresses resist material deformation. With perfect plasticity, the interface continues to deform unrealistically, and because the fine scales of the roll-up become unresolvable on the grid, leads to unphysical mixing in that region. By contrast, the simulation with strain hardening reaches a steady state shape within the simulation time.

To demonstrate where artificial diffusion is active in this problem, several artificial diffusivities are plotted in Figure 6.25 for the case with strain hardening, at $t^* = 1.875$. This is just after re-shock, which corresponds to the second panel in Figure 6.24. Both the artificial bulk viscosity and shear viscosity are active around the shock, as shown in Figure 6.25(a,b), but the bulk viscosity is much larger in magnitude, as expected. In addition, these artificial viscosities are also active around elastic shear waves, most clearly seen in Figure 6.25(b). The artificial diffusivities for the kinematic tensors, shown in Figure 6.25(c,d), are concentrated at the material interface (around $x^* = 3.25$), which is beginning to roll up due to interfacial shear. This localization is due to the switching function which focuses on regions of vorticity, not dilatation.

Finally, the effect of grid refinement is shown in Figure 6.26, which focuses on large-scale quantities, and Figure 6.27, which shows metrics that are strongly affected by small scales. The large-scale quantities include the spike and bubble locations, and the mixing width δ_w, defined as

$$\delta_w = \frac{1}{L_y} \int_x \int_y 4Y_1(1 - Y_1)dxdy \tag{6.73}$$

where Y_1 is the mass fraction of copper. The quantities sensitive to small scales include the equivalent plastic strain and enstrophy, defined as

$$\Omega = (\nabla \times \boldsymbol{u}) \cdot (\nabla \times \boldsymbol{u}) \tag{6.74}$$

In Figure 6.26, it is clear that strain hardening leads to plateaus in the mixing width and
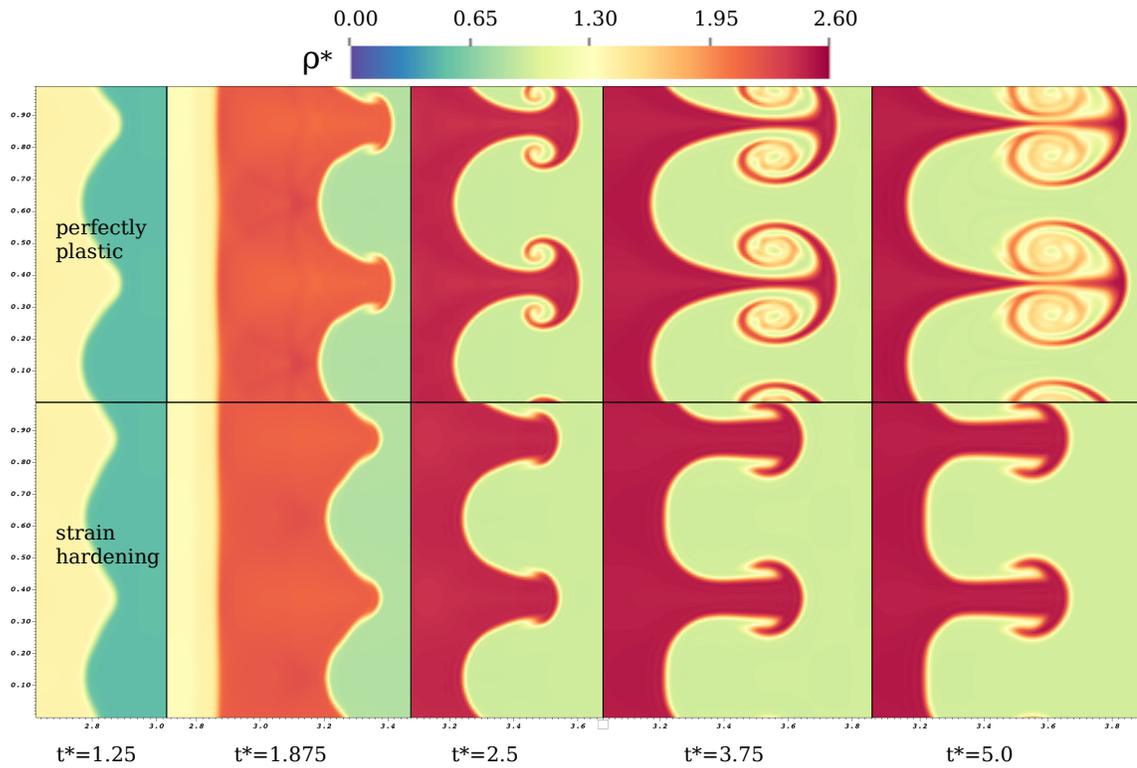
FIGURE 6.24: Snapshots of normalized density $\rho^* = \rho/\rho_0$ in a Richtmyer-Meshkov instability problem with $N_x = 768$ and $N_y = 128$, comparing perfect plasticity (top row) and strain hardening (bottom row). A moving sub-domain centered around the material interface is plotted. Time has been non-dimensionalized by the speed of sound in the copper material and the domain width L_y, and the spatial axes have been non-dimensionalized by L_y. The panel for $t^* = 1.875$ shows the development just after re-shock, and the left-traveling shock can be seen at $x^* \approx 2.9$.

(a)

(b)

(c)

(d)
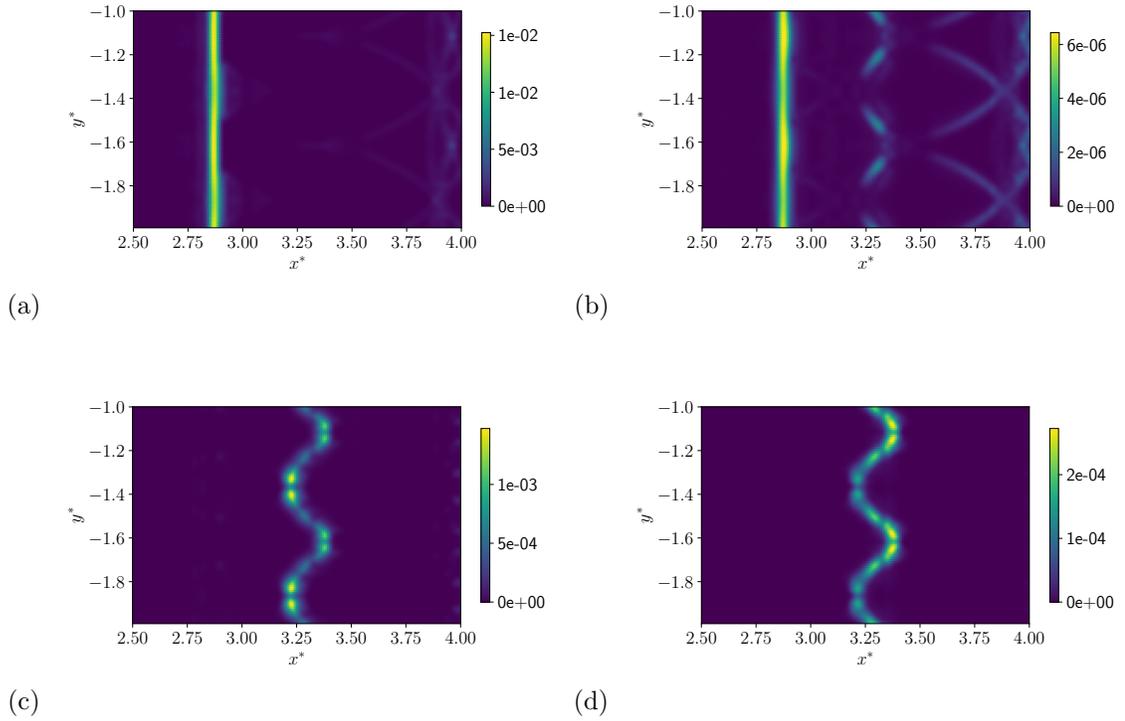
FIGURE 6.25: Illustration of where LAD coefficients are active in the Richtmyer-Meshkov problem with strain hardening at $t^* = 1.875$, just after re-shock. The shock is at $x^* \approx 2.9$, and the interface is at $x^* \approx 3.25$. All diffusivities are non-dimensionalized using the pressure scale $p_{\infty,1}$, the length scale L_y, and the density scale $\rho_{0,1}$ (a) Artificial bulk viscosity β^*, (b) artificial shear viscosity μ^*, (c) artificial diffusivity for $\underline{\mathbf{g}}^e$, g^{e*}, and (d) artificial diffusivity for $\underline{\mathbf{G}}^p$, g^{p*}. Artificial diffusivities for the kinematic tensors have been summed according to a volume fraction weighting between the two components.
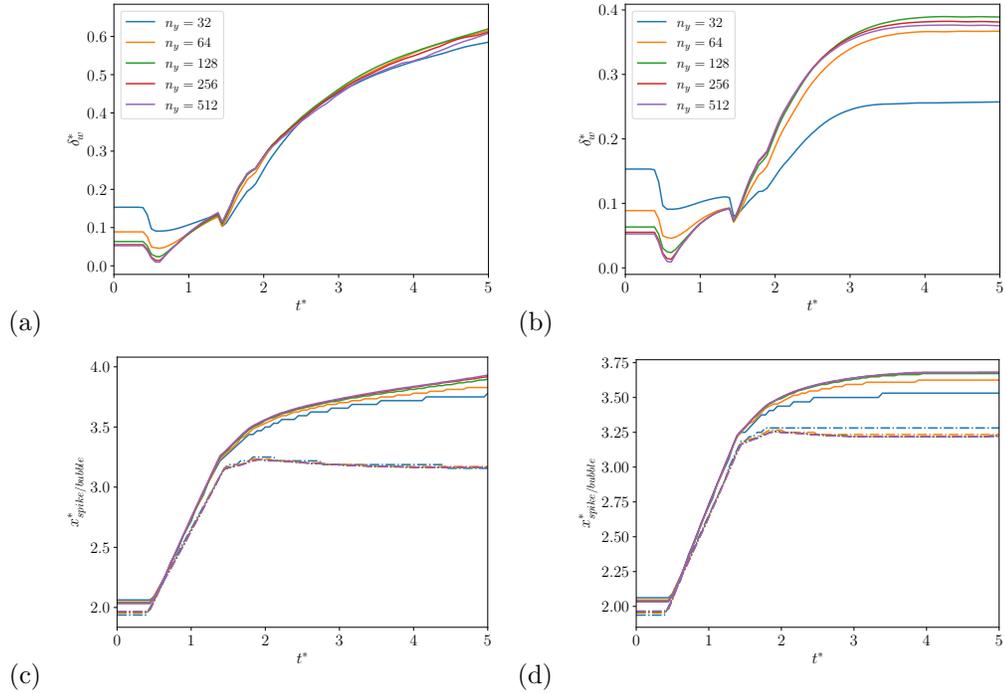
FIGURE 6.26: Quantitative metrics showing convergence behavior of large-scale features: (a) mixing width with perfect plasticity, (b) mixing width with strain hardening, (c) bubble and spike locations with perfect plasticity, and (d) bubble and spike locations with strain hardening. In plots (c) and (d), solid lines indicate spike locations and dash-dot lines indicate bubble positions.

spike/bubble locations, but with perfect plasticity, these quantities continue to evolve. Regardless of plasticity model, however, these larger scale quantities converge with grid refinement. Certain quantities do not converge with the grid refinement used here, and these are shown in Figure 6.27. As expected, the greater spike/bubble penetration and interface roll-up with perfect plasticity leads to greater plastic strain and enstrophy, and more is resolved as the grid is refined, indicating that these contributions come from large values of plastic strain localized to the fine scales in the simulation. At the finest grid resolution with strain hardening, the result is beginning to converge. In the enstrophy, spikes are observed when the shock hits the interface initially, and then at subsequent re-shocks. With strain hardening, the highest resolution case is beginning to converge the later time behavior, but the peaks do not converge because the baroclinic torque generated at the interface increases as both the shock and material interface become better refined.

From the standpoint of practical simulations, in which the grid resolution may be quite coarse, the large scale features can still be accurately predicted even though the quantities sensitive to small scales continue to change with grid refinement.
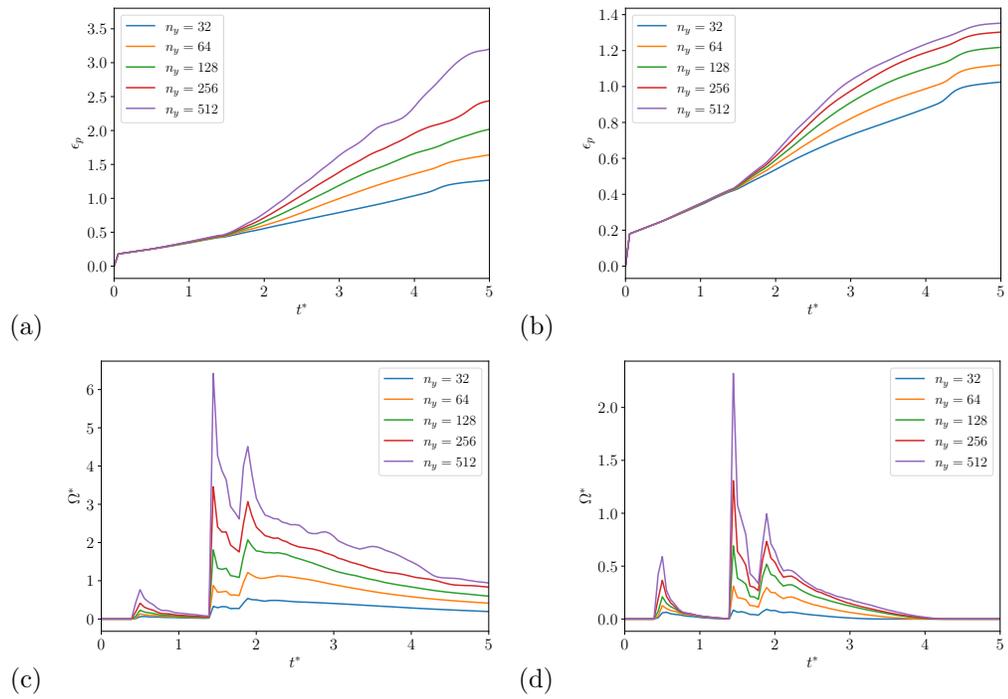
FIGURE 6.27: Quantitative metrics showing convergence behavior of features sensitive to small scales: (a) integrated equivalent plastic strain with perfect plasticity, (b) integrated equivalent plastic strain with strain hardening, (c) integrated enstrophy with perfect plasticity, and (d) integrated enstrophy with strain hardening.

6.6   Conclusion

In this work we have developed a high-order Eulerian finite difference method capable of simulating shocks in elastic-plastic media, as well as material interfaces between solids, and between solids and fluids. This method uses localized artificial diffusivity (LAD) to provide regularization around shocks and stabilize the governing equations. This is similar to previous efforts using LAD for gasdynamics and turbulence, but new LAD terms are added to the equations for kinematic tensors, $\underline{\mathbf{g}}^e$ and $\underline{\mathbf{G}}^p$, which are needed to compute elastic stresses, elastic energy, and plastic strain. This method is capable of simulating problems involving a variety of inelastic phenomena, including rate-independent and rate-dependent strain hardening. Strain hardening is emphasized here, but other phenomena, such as thermal softening, can be included if they can be described in terms of functional dependence of the shear modulus and yield stress on flow quantities, such as temperature, pressure, and plastic strain. The method has been demonstrated on a variety of test problems, including 1D impacts with and without strain hardening; a problem like the Shu-Osher problem, which demonstrates the capability of the method to resolve shocks interacting with other types of flow features; a Taylor impact problem, which demonstrates the ability to use real plasticity models and equations of state, and obtain results commensurate with Lagrangian methods in the literature; and finally a Richtmyer-Meshkov instability (RMI) between two solids, with and without strain hardening, which would be challenging to simulate using traditional Lagrangian methods due to mesh deformation. In the RMI problem, strain hardening and grid refinement are explored, showing how strain hardening arrests the development of the instability, and that large-scale features such as spike/bubble penetration converge rapidly, even as smaller scale features do not.

For future work, improvements to the method can still be made by increasing the order of accuracy of the implicit plastic relaxation. Another area for improvement is in the treatment of rotation, which limits to temporal first order. If the plasticity temporal error is reduced to higher order, then this error should be addressed as well to achieve that benefit. It would also be straightforward to use a finite plastic relaxation time, which would allow the simulation of a wider variety of plastic waves than simulated here, provided a suitable model for the relaxation timescale is available. Additional work can be done to improve the mixture equation of state, either by increasing the robustness of the pressure-temperature equilibrium algorithm, or by augmenting the current method to a 5 or 6 equation model, using the same kinematic equations and method as laid out here.

6.7   Appendix A: Governing equations in index notation

While most of the equations discussed in this work have been presented in tensor notation for ease of manipulation, we include here the governing equations in index notation for maximum clarity.

Conservation equations for mass, momentum and energy, corresponding to Equations 6.1-6.3:

$$\frac{\partial \rho Y_m}{\partial t} + \frac{\partial u_k \rho Y_m}{\partial x_k} = -\frac{\partial \left(J_m^*\right)_k}{\partial x_k}, \tag{6.75}$$

$$\frac{\partial \rho u_i}{\partial t} + \frac{\partial}{\partial x_k}\left[u_k \rho u_i - \sigma_{ik}\right] = \frac{\partial \tau_{ik}^*}{\partial x_k}, \tag{6.76}$$

$$\frac{\partial}{\partial t}\left[\rho\left(e + \frac{1}{2}u_i u_i\right)\right] + \frac{\partial}{\partial x_k}\left[u_k \rho\left(e + \frac{1}{2}u_i u_i\right) - u_i \sigma_{ik}\right] = \frac{\partial}{\partial x_k}\left[u_i \tau_{ik}^* - q_k^*\right]. \tag{6.77}$$

Governing equations for kinematic tensors without numerical regularization terms, corresponding to Equations 6.5-6.12:

$$\underbrace{\frac{\partial g_{ij}}{\partial t}}_{\substack{\text{local}\\\text{derivative}}} + \underbrace{u_k \frac{\partial g_{ij}}{\partial x_k}}_{\text{advection}} + \underbrace{g_{ik}\frac{\partial u_k}{\partial x_j}}_{\text{strain source}} = 0 \iff \underbrace{\frac{\partial g_{ij}}{\partial t}}_{\substack{\text{local}\\\text{derivative}}} + \underbrace{\frac{\partial g_{ik}u_k}{\partial x_j}}_{\substack{\text{curl-free}\\\text{advection/strain}}} + \underbrace{u_k\left(\frac{\partial g_{ij}}{\partial x_k} - \frac{\partial g_{ik}}{\partial x_j}\right)}_{\substack{\text{non-zero curl}\\\text{advection/strain}}} = 0, \tag{6.78}$$

$$\underbrace{\frac{\partial g_{ij}}{\partial t}}_{\substack{\text{local}\\\text{derivative}}} + \underbrace{\frac{\partial g_{ik}u_k}{\partial x_j}}_{\substack{\text{curl-free}\\\text{advection/strain}}} + \underbrace{u_k\left(\frac{\partial g_{ij}}{\partial x_k} - \frac{\partial g_{ik}}{\partial x_j}\right)}_{\substack{\text{non-zero curl}\\\text{advection/strain}}} = 0, \tag{6.79}$$

$$\underbrace{\frac{\partial g_{ij}^e}{\partial t}}_{\substack{\text{local}\\\text{derivative}}} + \underbrace{\frac{\partial g_{ik}^e u_k}{\partial x_j}}_{\substack{\text{curl-free}\\\text{advection/strain}}} + \underbrace{u_k\left(\frac{\partial g_{ij}^e}{\partial x_k} - \frac{\partial g_{ik}^e}{\partial x_j}\right)}_{\substack{\text{non-zero curl}\\\text{advection/strain}}} - \underbrace{L_{ik}^p g_{kj}^e}_{\text{elastic-plastic source}} = 0, \tag{6.80}$$

$$\underbrace{\frac{\partial g_{ij}^p}{\partial t}}_{\substack{\text{local}\\\text{derivative}}} + \underbrace{u_k \frac{\partial g_{ij}^p}{\partial x_k}}_{\text{advection}} + \underbrace{g_{ik}^p L_{kj}^p}_{\text{elastic-plastic source}} = 0, \tag{6.81}$$

$$\underbrace{\frac{\partial G_{ij}}{\partial t}}_{\substack{\text{local}\\\text{derivative}}} + \underbrace{u_k \frac{\partial G_{ij}}{\partial x_k}}_{\text{advection}} + \underbrace{\left(G_{ik}\frac{\partial u_k}{\partial x_j} + G_{jk}\frac{\partial u_k}{\partial x_i}\right)}_{\text{strain source}} = 0, \tag{6.82}$$

$$\underbrace{\frac{\partial G_{ij}^e}{\partial t}}_{\substack{\text{local}\\\text{derivative}}} + \underbrace{u_k \frac{\partial G_{ij}^e}{\partial x_k}}_{\text{advection}} + \underbrace{\left(G_{ik}^e\frac{\partial u_k}{\partial x_j} + G_{jk}^e\frac{\partial u_k}{\partial x_i}\right)}_{\text{strain source}} - \underbrace{g_{ki}^e\left(L_{kl}^p + L_{lk}^p\right)g_{lj}^e}_{\text{elastic-plastic source}} = 0, \tag{6.83}$$

$$\underbrace{\frac{\partial G_{ij}^p}{\partial t}}_{\substack{\text{local}\\\text{derivative}}} + \underbrace{u_k \frac{\partial G_{ij}^p}{\partial x_k}}_{\text{advection}} + \underbrace{L_{ki}^p G_{kj}^p + G_{ik}^p L_{kj}^p}_{\text{elastic-plastic source}} = 0, \tag{6.84}$$

$$\frac{\partial \rho S^p}{\partial t} + \frac{\partial \rho S^p u_i}{\partial x_i} = \frac{1}{T}\left[L^{S^p} + \tau_{ij}^*\frac{\partial u_j}{\partial x_i} - \frac{\partial q_i^*}{\partial x_i}\right]. \tag{6.85}$$

Governing equations for kinematic tensors $\underline{\mathbf{g}}^e$ and $\underline{\mathbf{G}}^p$ for a multi-material mixture, with numerical regularization terms included:

$$\underbrace{\frac{\partial (g_m^e)_{ij}}{\partial t}}_{\substack{\text{local} \\ \text{derivative}}} + \underbrace{\frac{\partial (g_m^e)_{ik} u_k}{\partial x_j}}_{\substack{\text{curl-free} \\ \text{advection/strain}}} + \underbrace{u_k \left(\frac{\partial (g_m^e)_{ij}}{\partial x_k} - \frac{\partial (g_m^e)_{ik}}{\partial x_j} \right)}_{\substack{\text{non-zero curl} \\ \text{advection/strain}}} - \underbrace{(L_m^p)_{ik} (g_m^e)_{kj}}_{\text{elastic-plastic source}} =$$

$$\underbrace{\frac{\zeta}{\Delta t} \overline{\left(\frac{\rho}{\rho_0 \, |\boldsymbol{g^e}|} - 1 \right)} (g_m^e)_{ij}}_{\text{determinant compatibility}} + \underbrace{\frac{\partial}{\partial x_k} \left(g^{e*} \frac{\partial (g_m^e)_{ij}}{\partial x_k} \right)}_{\text{artificial diffusion}} \quad (6.86)$$

$$\underbrace{\frac{\partial (G_m^p)_{ij}}{\partial t}}_{\substack{\text{local} \\ \text{derivative}}} + \underbrace{u_k \frac{\partial (G_m^p)_{ij}}{\partial x_k}}_{\text{advection}} + \underbrace{(L_m^p)_{ki} (G_m^p)_{kj} + (G_m^p)_{ik} (L_m^p)_{kj}}_{\text{elastic-plastic source}} = \underbrace{\frac{\zeta}{\Delta t} \overline{\left(\frac{1}{|\boldsymbol{G^p}|} - 1 \right)} (G_m^p)_{ij}}_{\text{determinant compatibility}} + \underbrace{\frac{\partial}{\partial x_k} \left(g^{p*} \frac{\partial (G_m^p)_{ij}}{\partial x_k} \right)}_{\text{artificial diffusion}}$$

$$(6.87)$$

Chapter 7

Assessment of diffuse-interface methods for compressible multiphase fluid flows and elastic-plastic deformation in solids

The contents of this chapter have been published in the Journal of Computational Physics [98].

7.1 Introduction

Compressible multiphase fluid flow and multiphase elastic-plastic deformation of solid materials with strength are important phenomena in many engineering applications, including shock compression of condensed matter, detonations and shock-material-interface interactions, impact welding, high-speed fuel atomization and combustion, and cavitation and bubble collapse motivated by both mechanical and biomedical systems. In this work, we are concerned with the numerical modeling of multiphase systems, i.e., those systems that involve two or more phases of gas, liquid, or solid in the domain. The numerical simulation of these multiphase systems presents several new challenges in addition to those associated with analogous single-phase simulations. These modeling complications include but are not limited to (1) representing the phase interface on an Eulerian grid; (2) resolving discontinuities in quantities at the interface, especially for high-density ratios; (3) maintaining conservation of (a) the mass of each phase, (b) the mixture momentum, and (c) the total energy of the system; and (4) achieving an accurate mixture representation of the interface for maintaining thermodynamic equilibria. Hence, the numerical modeling of multiphase compressible fluid flows and deformation

of solid materials with strength are still an active area of research.

With these numerical challenges in mind, we choose to pursue the single-fluid approach [106], in which a single set of equations is solved to describe all of the phases in the domain, as opposed to a multi-fluid approach, which requires solving a separate set of equations for each of the phases. We are presented with various choices in terms of the system of equations that can be used to represent a compressible multiphase system. In this work, we employ a multicomponent system of equations (a four-equation model) that assumes spatially local pressure and temperature equilibria between phases, including at locations within the diffuse material interface [181, 202, 125, 40]. The underlying assumption is that the temperature between two phases reaches equilibrium faster than any other time scale in the flow. Relaxing the assumption of temperature equilibrium, Allaire et al. [6] and Kapila et al. [105] developed the five-equation model that has proven successful for a variety of applications with high density ratios, strong compressibility effects, and phases with disparate equations of state (EOS), and has been widely adopted for the simulation of compressible two-phase flows [180, 184, 8, 179, 42, 196, 160, 210, 37, 69, 70, 95, 97, 93]. Finally, there are six- and seven-equation models that are more general and include more non-equilibrium effects but are not as widely used for the simulation of two-phase flows [213, 11, 172, 174].

For representing the interface on an Eulerian grid, we use an interface-capturing method, as opposed to an interface-tracking method, due to the natural ability of the former method to simulate dynamic creation of interfaces and topological changes [136]. Interface-capturing methods can be classified into sharp-interface and diffuse-interface methods. In this work, we choose to use diffuse-interface methods for modeling the interface between compressible materials [175]. This choice is due to the natural advantages that the diffuse-interface methods offer over the sharp-interface methods, such as ease of representation of the interface, low cost, good conservation properties, and parallel scalability.

Historically, diffuse-interface methods for compressible flows involved modeling the interface implicitly, i.e., with no explicit interface capturing through regularization/reconstruction. These methods can be classified as implicit diffuse-interface methods. These methods assume that the underlying numerical methods are capable of handling the material interfaces, a concept similar to implicit large-eddy simulation. One challenge with the implicit diffuse-interface capturing of material interfaces is that the interface tends to diffuse over time. Unlike shock waves, in which the convective characteristics sharpen the shock over time, material interfaces (like contact discontinuities) do not sharpen naturally; therefore, modeling material interfaces requires an active balance between interface sharpening and diffusion to maintain an appropriate interface thickness over time. Therefore, in the present work, the focus is on explicit diffuse-interface methods. These methods explicitly model the interface using the interface regularization/reconstruction techniques.

This paper explores three explicit diffuse-interface methods that are representative of the different approaches for this problem and possesses unique characteristics. The first approach (referred to as

the LAD approach) is based on the localized-artificial-diffusivity (LAD) method [39, 188, 5], in which localized, nonlinear diffusion terms are added to the individual phase mass transport equations and coupled with the other conservation equations. This method conserves the mass of individual phases, mixture momentum, and total energy of the system due to the conservative nature of the diffusion terms added to the system of equations and results in no net mixture-mass transport. This method is primarily motivated by applications involving miscible, multicomponent, single-phase flows, but it has been successfully adapted for multiphase flow applications. The idea behind this approach is to effectively add species diffusion in the selected regions of the domain to properly resolve the interface on the grid and to prevent oscillations due to discontinuities in the phase mass equations. High-order compact derivative schemes can be used to discretize the added diffusion terms without resulting in distortion of the shape of the interface over long-duration time advancement. However, one drawback of this approach is that the interface thickness increases with time due to the lack of sharpening fluxes that act against the diffusion. This method is therefore most effective for problems in which the interface is in compression (such as shock/material-interface interactions with normal alignment). However, the deficiency of this method due to the lack of a sharpening term is evident for applications in which the interface between immiscible materials undergoes shear or expansion/tension. LAD formulations have also been examined in the context of five-equation models, in which localized diffusion is also added to the volume fraction transport equation [9].

The second approach (referred to as the gradient-form approach) is based on the quasi-conservative method proposed by Shukla et al. [180], in which diffusion and sharpening terms (together called regularization terms) are added for the individual phase volume fraction transport equations and coupled with the other conservation equations [196]. This method only approximately conserves the mass of individual phases, mixture momentum, and total energy of the system due to the non-conservative nature of the regularization terms added to the system of equations. In contrast to the LAD approach, this method can result in net mixture-mass transport, which can sharpen or diffuse the mixture density; depending on the application, this may be an advantageous or disadvantageous property. The primary advantage of this method is that the regularization terms are insensitive to the method of discretization; they can be discretized using high-order compact derivative schemes without distorting the shape of the interface over long-duration time advancement. However, the non-conservative nature of this approach results in poor performance of the method for certain applications. For example, premature topological changes and unphysical interface behavior can be observed when the interfaces are poorly resolved (exacerbating the conservation error) and subjected to shocks that are not aligned with the interface.

The third approach (referred to as the divergence-form approach) is based on the fully conservative method proposed by Jain et al. [97], in which diffusion and sharpening terms are added to the individual phase volume fraction transport equations and coupled with the other conservation equations. This method conserves the mass of individual phases, mixture momentum, and total

energy of the system due to the conservative nature of the regularization terms added to the system of equations. Similar to the gradient-form approach and in contrast to the LAD approach, this method can result in net mixture-mass transport, which can sharpen or diffuse the mixture density. The primary challenge of this method is that one needs to be careful with the choice of discretization used for the regularization terms. Using a second-order finite-volume scheme (in which the nonlinear fluxes are formed on the faces), Jain et al. [97] showed that a discrete balance between the diffusion and sharpening terms is achieved, thereby eliminating the spurious behavior that was discussed by Shukla et al. [180]. The idea behind this is similar to the use of the balanced-force algorithm [63, 131] for the implementation of the surface-tension forces, in which a discrete balance between the pressure and surface-tension forces is necessary to eliminate the spurious velocity around the interface. The current study also demonstrates that appropriately crafted higher-order schemes may be used to effectively discretize the regularization terms. This method is free of premature topological changes and unphysical interface behavior present with the previous approach. However, due to the method of discretization, the anisotropy of the derivative scheme can more significantly distort the shape of the material interface over long-duration time advancement in comparison to the gradient-form approach; the severity of this problem is significantly reduced when using higher-order schemes.

For all the three diffuse-interface methods considered in this work, it is important to include physically consistent corrections, associated with the interface regularization process, in each of the governing equations. For example, Cook [40], Tiwari et al. [196], and Jain et al. [97] discuss physically consistent regularization terms for the LAD, gradient-form, and divergence-form approaches, respectively. The physically consistent regularization terms of Cook [40], Tiwari et al. [196], and Jain et al. [97] are derived in such a way that the regularization terms do not spuriously contribute to the kinetic energy and entropy of the system. This significantly improves the stability of the simulation, especially for flows with high density ratios. However, discrete conservation of kinetic energy and entropy is needed to show the stability of the methods for high-Reynolds-number turbulent flows [94].

We employ a fully Eulerian method for modeling the deformation of solid materials, as opposed to a fully Lagrangian approach [26] or a mixed approach such as arbitrary-Lagrangian-Eulerian methods [52], because of its cost-effectiveness and accuracy to handle large deformations. There are various Eulerian approaches in the literature that differ in the way the deformation of the material is tracked. The popular methods employ the inverse deformation gradient tensor [135, 148, 73], the left Cauchy-Green tensor [189, 190], the co-basis vectors [59], the initial material location [200, 96], or other variants of these methods to track the deformation of the material in the simulation. In this work, we use the inverse deformation gradient tensor approach because of its applicability to model plasticity. We propose consistent corrections in the kinematic equations, that describe the deformation of the solid, associated with the interface regularization process.

In summary, the two main objectives of this paper are as follows. The first objective is to

assess several diffuse-interface-capturing methods for compressible two-phase flows. The interface-capturing methods in this work will be used with a four-equation multicomponent model; however, they are readily compatible with a variety of other models, including the common five-, six-, or seven-equation models. The second objective is to extend these interface-capturing methods for the simulation of elastic-plastic deformation in solid materials with strength, including comparison of these methods in the context of modeling interfaces between solid materials.

The remainder of this paper is outlined as follows: Section 7.2 describes the three diffuse-interface methods considered in this study, along with the details of their implementation. Section 7.3 discusses the application of these methods to a variety of problems including a shock/helium-bubble interaction in air, an advecting air bubble in water, a shock/air-bubble interaction in water, and a Richtmyer–Meshkov instability of an interface between copper and aluminum. Concluding remarks are made in Section 7.4 along with a summary. A table highlighting the strengths and limitations of the different methods considered in this work is also presented in this section.

7.2 Theoretical and numerical model

7.2.1 Governing equations

The governing equations for the evolution of the multiphase flow or multimaterial continuum in conservative Eulerian form are described in Eqs. (7.1)-(7.3). This consists of the conservation of species mass (Eq. 7.1), total momentum (Eq. 7.2), and total energy (Eq. 7.3). These are followed by the kinematic equations that track the material deformation, which include transport equations for the elastic component of the inverse deformation gradient tensor (Eq. 7.4), and the plastic Finger tensor (Eq. 7.5).

$$\underbrace{\frac{\partial \rho Y_m}{\partial t}}_{\substack{\text{local} \\ \text{derivative}}} + \underbrace{\frac{\partial u_k \rho Y_m}{\partial x_k}}_{\text{advection}} = - \underbrace{\frac{\partial \left(J_m^* \right)_i}{\partial x_i}}_{\substack{\text{artificial} \\ \text{diffusion}}} + \underbrace{J_m}_{\substack{\text{interface} \\ \text{regularization}}}, \tag{7.1}$$

$$\underbrace{\frac{\partial \rho u_i}{\partial t}}_{\substack{\text{local} \\ \text{derivative}}} + \frac{\partial}{\partial x_k} \left(\underbrace{u_k \rho u_i}_{\text{advection}} - \underbrace{\sigma_{ik}}_{\substack{\text{stress} \\ \text{source}}} \right) = \underbrace{\frac{\partial \tau_{ik}^*}{\partial x_k}}_{\substack{\text{artificial} \\ \text{diffusion}}} + \underbrace{F_i}_{\substack{\text{interface} \\ \text{regularization}}}, \tag{7.2}$$

$$\underbrace{\frac{\partial}{\partial t} \left[\rho \left(e + \frac{1}{2} u_j u_j \right) \right]}_{\substack{\text{local} \\ \text{derivative}}} + \frac{\partial}{\partial x_k} \left[\underbrace{u_k \rho \left(e + \frac{1}{2} u_j u_j \right)}_{\text{advection}} - \underbrace{u_i \sigma_{ik}}_{\substack{\text{stress} \\ \text{source}}} \right]$$

$$= \underbrace{\frac{\partial}{\partial x_k} \left(u_i \tau_{ik}^* - q_k^* \right)}_{\text{artificial diffusion}} + \underbrace{H}_{\substack{\text{interface} \\ \text{regularization}}}, \tag{7.3}$$

$$
\underbrace{\frac{\partial g_{ij}^e}{\partial t}}_{\substack{\text{local} \\ \text{derivative}}} + \underbrace{\frac{\partial g_{ik}^e u_k}{\partial x_j}}_{\substack{\text{curl-free} \\ \text{advection/strain}}} + \underbrace{u_k \left(\frac{\partial g_{ij}^e}{\partial x_k} - \frac{\partial g_{ik}^e}{\partial x_j} \right)}_{\substack{\text{non-zero curl} \\ \text{advection/strain}}} - \underbrace{\frac{1}{2\mu\tau_{rel}} g_{ik}^e \sigma_{kj}'}_{\text{elastic-plastic source}}
$$

$$
= \underbrace{\frac{\zeta^e}{\Delta t} \left(\frac{\rho}{\rho_0 \det(\boldsymbol{g^e})} - 1 \right) g_{ij}^e}_{\text{density compatibility}} + \underbrace{\frac{\partial}{\partial x_k} \left(g^{e*} \frac{\partial g_{ij}^e}{\partial x_k} \right)}_{\text{artificial diffusion}} + \underbrace{K_{ij}^e}_{\substack{\text{interface} \\ \text{regularization}}} ,
\tag{7.4}
$$

$$
\underbrace{\frac{\partial G_{ij}^p}{\partial t}}_{\substack{\text{local} \\ \text{derivative}}} + \underbrace{u_k \frac{\partial G_{ij}^p}{\partial x_k}}_{\text{advection}} + \underbrace{\frac{1}{2\mu\tau_{rel}} \left(G_{ik}^p g_{kl}^e \sigma_{lm}' (g^e)_{mj}^{-1} + G_{jk}^p g_{kl}^e \sigma_{lm}' (g^e)_{mi}^{-1} \right)}_{\text{elasto-plastic source}}
$$

$$
= \underbrace{\frac{\zeta^p}{\Delta t} \left(\frac{1}{\det(\boldsymbol{G^p})^{1/2}} - 1 \right) G_{ij}^p}_{\text{density compatibility}} + \underbrace{\frac{\partial}{\partial x_k} \left(g^{p*} \frac{\partial G_{ij}^p}{\partial x_k} \right)}_{\text{artificial diffusion}} .
\tag{7.5}
$$

Here, t and $\boldsymbol{x}$ represent time and the Eulerian position vector, respectively. Y_m describes the mass fraction of each constituent material, m. The variables $\boldsymbol{u}$, ρ, e, and $\underline{\boldsymbol{\sigma}}$ describe the mixture velocity, density, internal energy, and Cauchy stress, respectively, which are related to the species-specific components by the relations $\rho = \sum_{m=1}^M \phi_m \rho_m$, $e = \sum_{m=1}^M Y_m e_m$, and $\underline{\boldsymbol{\sigma}} = \sum_{m=1}^M \phi_m \underline{\boldsymbol{\sigma}}_m$, in which ϕ_m is the volume fraction of material m, and M is the total number of material constituents. The variables g_{ij}^e and G_{ij}^p are tensors that track elastic and plastic material deformation in problems with solids. These equations are described in greater detail in the next section.

The right-hand-side terms describe the localized artificial diffusion (see also Section 7.2.5), including the artificial viscous stress, $\tau_{ik}^* = 2\mu^* S_{ik} + (\beta^* - 2\mu^*/3) (\partial u_j / \partial x_j) \delta_{ik}$, and the artificial enthalpy flux, $q_i^* = -\kappa^* \partial T / \partial x_i + \sum_{m=1}^M h_m (J_m^*)_i$, with strain rate tensor, $S_{ik} = (\partial u_i / \partial x_k + \partial u_k / \partial x_i) /2$, and temperature, T. The second term in the artificial enthalpy flux expression is the enthalpy diffusion term [40], in which $h_m = e_m + p_m / \rho_m$ is the enthalpy of species m. The artificial Fickian diffusion of species m is described by $(J_m^*)_i = -\rho [D_m^* (\partial Y_m / \partial x_i) - Y_m \sum_k D_k^* (\partial Y_k / \partial x_i)]$.

7.2.2 Material deformation and plasticity model

The kinematic equations that describe the deformation of the solid in the Eulerian framework employ the inverse deformation gradient tensor, $g_{ij} = \partial X_i / \partial x_j$, in which $\boldsymbol{X}$ and $\boldsymbol{x}$ describe the position of a continuum parcel in the material (Lagrangian) and spatial (Eulerian) perspectives, respectively. In this work, a single inverse deformation gradient is used to describe the kinematics of the mixture [72, 73]. Following [135], a multiplicative decomposition of the total inverse deformation gradient tensor, $\boldsymbol{g}$, into elastic, $\boldsymbol{g^e}$, and plastic, $\boldsymbol{g^p}$, components is assumed, $g_{ij} = g_{ik}^p g_{kj}^e$, reflecting the assumption that the plastic deformation is recovered when the elastic deformation is reversed, $g_{ij}^p = g_{ik} (g^e)_{kj}^{-1}$. It is additionally assumed that the plastic deformation is volume preserving [162], providing compatibility conditions for the inverse deformation gradient tensor determinants,

$\det\left(\boldsymbol{g^p}\right) = 1$ and $\det\left(\boldsymbol{g}\right) = \det\left(\boldsymbol{g^e}\right) = \rho/\rho_0$, in which ρ_0 represents the undeformed density and $\det\left(\cdot\right)$ represents the determinant operator. In this work, the plastic Finger tensor $G_{ij}^p = g_{ik}^p g_{jk}^p$ is solved for because it tends to be more stable than the equation for $\mathbf{g}^p$, and because models for strain hardening are often parametrized in terms of norms of the plastic Finger tensor. This choice and its alternatives are discussed in detail in [206].

We also assume that the materials with strength are elastic perfectly plastic, i.e., the material yield stress is independent of strain and strain rate; thus, only the elastic component of the inverse deformation gradient tensor is necessary to close the governing equations. As a result, we solve only the equation for elastic deformation in the present work. The plastic component of the inverse deformation gradient tensor, or the full tensor, can be employed to supply the plastic strain and strain rate necessary for more general plasticity models [5].

Plastic deformation is incorporated into the numerical framework by means of a visco-elastic Maxwell relaxation model, which has been employed recently in several Eulerian approaches [144, 149, 73]. The plastic relaxation timescale is described by

$$\frac{1}{\tau_{\mathrm{rel}}} = \frac{1}{(\rho/\rho_0)\,\tau_0} \left[\frac{\mathrm{R}\left(\|\boldsymbol{\sigma}'\|^2 - \frac{2}{3}\sigma_Y^2\right)}{\mu^2}\right], \tag{7.6}$$

in which $\underline{\boldsymbol{\sigma}}' = \mathrm{dev}\left(\boldsymbol{\sigma}\right)$ and μ is the material shear modulus. The ramp function $R\left(x\right) = \max\left(x, 0\right)$ turns on plasticity effects only when the yield criterion is satisfied. In many cases, the elastic-plastic source term is stiff due to the small value of τ_{rel} relative to the convective deformation scales. To overcome this time step restriction, implicit plastic relaxation is performed at each timestep, based on the method of [59] and described by [73].

7.2.3 Equations of state and constitutive equations

A hyperelastic constitutive model, in which the elastic stress–strain relationship is compatible with a strain energy-density functional, is assumed to close the thermodynamic relationships in the governing equations. The internal energy, e, is additively decomposed into a hydrodynamic component, e_h, and an elastic component, e_e, as in [144]. The hydrodynamic component is analogous to a stiffened gas, with

$$e = e_h\left(p, \rho\right) + e_e\left(\hat{\boldsymbol{g}}\right), \qquad e_h = \frac{p + \gamma p_\infty}{(\gamma - 1)\rho}, \qquad e_e = \frac{\mu}{4\rho_0}\mathrm{tr}\left[\left(\hat{\boldsymbol{g}} - \boldsymbol{I}\right)^2\right], \tag{7.7}$$

in which $\hat{\underline{\boldsymbol{g}}} = \det\left(\underline{\boldsymbol{G^e}}\right)^{-1/3}\underline{\boldsymbol{G^e}}$, $\underline{\boldsymbol{G^e}} = \boldsymbol{g}^{eT}\boldsymbol{g^e}$, p is the pressure, p_∞ (with units of pressure) and γ (nondimensional) are material constants of the stiffened gas model for the hydrodynamic component of internal energy. With this EOS, the Cauchy stress, $\underline{\boldsymbol{\sigma}}$, satisfying the Clausius-Duhem inequality

is described by

$$\underline{\boldsymbol{\sigma}} = -p\underline{\boldsymbol{I}} - \mu\frac{\rho}{\rho_0}\left\{\det\left(\underline{\boldsymbol{G^e}}\right)^{-2/3}\mathrm{dev}\left[\left(\underline{\boldsymbol{G^e}}\right)^2\right] - \det\left(\underline{\boldsymbol{G^e}}\right)^{-1/3}\mathrm{dev}\left(\underline{\boldsymbol{G^e}}\right)\right\}, \tag{7.8}$$

in which $\mathrm{dev}\left(\underline{\boldsymbol{G^e}}\right)$ signifies the deviatoric component of the tensor: $\mathrm{dev}\left(\underline{\boldsymbol{G^e}}\right) = \underline{\boldsymbol{G^e}} - \frac{1}{3}\mathrm{tr}\left(\underline{\boldsymbol{G^e}}\right)\underline{\boldsymbol{1}}$, with $\mathrm{tr}\left(\cdot\right)$ signifying the trace of the tensor and $\underline{\boldsymbol{1}}$ signifying the identity tensor. The elastic component of the internal energy, e_e, is assumed to be isentropic. Therefore, the temperature, T, and entropy, η, are defined by the hydrodynamic stiffened gas component of the EOS, as follows.

$$e_h = C_v T\left(\frac{p + \gamma p_\infty}{p + p_\infty}\right), \qquad R = C_p - C_v, \qquad \gamma = \frac{C_p}{C_v},$$
$$\eta - \eta_0 = C_p \ln\left(\frac{T}{T_0}\right) + R\ln\left(\frac{p_0 + p_\infty}{p + p_\infty}\right). \tag{7.9}$$

Here, η_0 is the reference entropy at pressure, p_0, and temperature, T_0. In the case of compressible flow with no material strength, the material model reduces to the stiffened gas EOS commonly employed for liquid/gas-interface interactions [180, 97].

7.2.4 Pressure and temperature equilibration method

Many models for multiphase simulation assume that the thermodynamic variables are not in equilibrium, necessitating the solution of an additional equation for volume fraction transport [180, 97]. Our model begins with the assumption that both pressure and temperature remain in equilibrium between the phases. The equilibration method follows from Cook [40] and Subramaniam et al. [188]. For a mixture of M species, we solve for $2M + 2$ unknowns, including the equilibrium pressure (p), the equilibrium temperature (T), the component volume fractions (ϕ_m), and the component internal energies (e_m), from the following equations.

$$p = p_m, \qquad T = T_m, \qquad \sum_{m=1}^{M}\phi_m = 1, \qquad \sum_{m=1}^{M}Y_m e_m = e. \tag{7.10}$$

To achieve a stable equilibrium, it requires that all phases be present with non-negative volume fractions throughout the entire simulation domain. This is achieved by initializing the problem with a minimum volume fraction (typically $\phi_{\min} = 10^{-6} - 10^{-4}$) and including additional criteria for volume fraction diffusion (Sections 7.2.6 and 7.2.7) or mass fraction diffusion (Section 7.2.5) based on out-of-bounds values of volume fraction and/or mass fraction. This equilibration method is stable in the well-mixed interface region, but can result in stability issues outside of the interface region, where the volume fraction of a material tends to become very small—a phenomenon exacerbated by high-order discretization methods.

7.2.5 Localized artificial diffusivity

LAD methods have long proven useful in conjunction with high-order compact derivative schemes to provide necessary solution-adaptive and localized diffusion to capture discontinuities and introduce a mechanism for subgrid dissipation. Regardless of the choice of interface-capturing method, LAD is required in the momentum, energy, and kinematic equations, in all calculations, to provide necessary regularization. For instance, the artificial shear viscosity, μ^*, primarily serves as a subgrid dissipation model, whereas the artificial bulk viscosity, β^*, enables shock capturing, and the artificial thermal conductivity, κ^*, captures contact discontinuities. The artificial kinematic diffusivities (g^{e*} and g^{p*}) facilitate capturing of strain discontinuities, particularly in regions of sustained shearing.

When LAD is also used for interface regularization (to capture material interfaces), the artificial diffusivity of species m, D_m^*, is activated, in which the coefficient C_D controls the interface diffusivity and the coefficient C_Y controls the diffusivity when the mass fraction goes out of bounds. When using the volume-fraction-based approaches for interface regularization (Sections 7.2.6 and 7.2.7), it is often unnecessary to also include the species LAD ($D_m^* = 0$); however, the species LAD seems to be necessary for some problems in conjunction with these other interface regularization approaches.

The artificial diffusivities are described below, where the overbar denotes a truncated Gaussian filter applied along each grid direction; Δ_i is the grid spacing in the i direction; $\Delta_{i,\mu}$, $\Delta_{i,\beta}$, $\Delta_{i,\kappa}$, Δ_{i,Y_m}, and $\Delta_{i,g}$ are weighted grid length scales in direction i; c_s is the linear longitudinal wave (sound) speed; H is the Heaviside function; and $\varepsilon = 10^{-32}$.

$$\mu^* = C_\mu \rho \overline{\left| \sum_{k=1}^{3} \frac{\partial^r S}{\partial x_k^r} \Delta_k^r \Delta_{k,\mu}^2 \right|}; \qquad \Delta_{i,\mu} = \Delta_i. \tag{7.11}$$

$$\beta^* = C_\beta \rho f_{sw} \overline{\left| \sum_{k=1}^{3} \frac{\partial^r (\boldsymbol{\nabla} \cdot \boldsymbol{u})}{\partial x_k^r} \Delta_k^r \Delta_{k,\beta}^2 \right|}; \qquad \Delta_{i,\beta} = \Delta_i \frac{\left(\frac{\partial \rho}{\partial x_i} \right)^2}{\sum_{k=1}^{3} \left(\frac{\partial \rho}{\partial x_k} \right)^2 + \varepsilon}. \tag{7.12}$$

$$\kappa^* = C_\kappa \frac{\rho c_s}{T} \overline{\left| \sum_{k=1}^{3} \frac{\partial^r e_h}{\partial x_k^r} \Delta_k^r \Delta_{k,\kappa} \right|}; \qquad \Delta_{i,\kappa} = \Delta_i \frac{\left(\frac{\partial e_h}{\partial x_i} \right)^2}{\sum_{k=1}^{3} \left(\frac{\partial e_h}{\partial x_k} \right)^2 + \varepsilon}. \tag{7.13}$$

$$D_m^* = \max\left\{ C_D c_s \overline{\left| \sum_{k=1}^{3} \frac{\partial^r Y_m}{\partial x_k^r} \Delta_k^r \Delta_{k,D} \right|}, \; C_Y \overline{\frac{c_s}{2}\left(|Y_m| - 1 + |1 - Y_m|\right) \sum_{k=1}^{3} \Delta_{k,Y}} \right\};$$

$$\Delta_{i,D} = \Delta_i \frac{\left(\frac{\partial Y_m}{\partial x_i}\right)^2}{\sum_{k=1}^{3}\left(\frac{\partial Y_m}{\partial x_k}\right)^2 + \varepsilon}; \qquad \Delta_{i,Y} = \Delta_i \frac{\left|\frac{\partial Y_m}{\partial x_i}\right|}{\sqrt{\sum_{k=1}^{3}\left(\frac{\partial Y_m}{\partial x_k}\right)^2 + \varepsilon}}. \tag{7.14}$$

$$g^* = C_g c_s \overline{\left| \sum_{k=1}^{3} \frac{\partial^r E^g}{\partial x_k^r} \Delta_k^r \Delta_{k,g} \right|}; \qquad \Delta_{i,g} = \Delta_i \frac{\left(\frac{\partial E^g}{\partial x_i}\right)^2}{\sum_{k=1}^{3}\left(\frac{\partial E^g}{\partial x_k}\right)^2 + \varepsilon}. \tag{7.15}$$

$$f_{sw} = H\left(-\boldsymbol{\nabla} \cdot \boldsymbol{u}\right) \frac{\left(\boldsymbol{\nabla} \cdot \boldsymbol{u}\right)^2}{\left(\boldsymbol{\nabla} \cdot \boldsymbol{u}\right)^2 + |\boldsymbol{\nabla} \times \boldsymbol{u}|^2 + \varepsilon}. \tag{7.16}$$

Here, $S = \sqrt{S_{ij}S_{ij}}$ is a norm of the strain rate tensor, S, and $E^g = \sqrt{\frac{2}{3}E_{ij}^g E_{ji}^g}$ is a norm of the Almansi finite-strain tensor associated with the g equations, $E_{ij}^g = \frac{1}{2}\left(\delta_{ij} - g_{ki}g_{kj}\right)$. The artificial kinematic diffusivities, g^{e*} and g^{p*}, are obtained by using the equation for g^*, but with the Almansi strain based on only the elastic or plastic component of the inverse deformation gradient tensor, respectively. We observe that LAD is not strictly necessary to ensure stability for the $\mathbf{g}^e$ equations; in fact, it has not been included in previous simulations [73, 188], because the elastic deformation is often small relative to the plastic deformation, but LAD is necessary to provide stability for the $\mathbf{G}^p$ equations, especially when the interface is re-shocked, resulting in sharper gradients in the plastic deformation relative to the elastic deformation.

Typical values for the model coefficients are $\zeta^e = \zeta^p = 0.5$, $C_\mu = 2 \times 10^{-3}$, $C_\beta = 1$, $C_\kappa = 1 \times 10^{-2}$, $C_D = 3 \times 10^{-3}$, $C_Y = 1 \times 10^2$, and $C_g = 1$; these values are used in the subsequent simulations unless stated otherwise. However, these coefficients often need to be specifically tailored to the problem; for example, the bulk viscosity coefficient can be increased to more effectively capture strong shocks in materials with large stiffening pressures. The procedure for tuning these coefficients is to use a 1D stationary discontinuity problem to find an acceptable compromise between reducing Gibbs phenomena and avoiding excessive discontinuity spreading. A rule of thumb to achieve this is to reduce the amplitude of Gibbs phenomena to 0.5% of the jump. An example of this tuning procedure for shocks and contact discontinuities in air, water, and steel is provided by Subramaniam et al. [188]. However, these coefficients do not need to be tailored to the interface-regularization method. For the results shown in Section 7.3, the coefficients only vary with the test problem and are kept the same when varying the method, in order to facilitate a fair comparison.

7.2.6 Fully conservative divergence-form approach to interface regularization

In this method, interface regularization is achieved with the use of diffusion and sharpening terms that balance each other. This results in constant interface thickness during the simulation, unlike the LAD method, in which the interface thickness increases over time due to the absence of interface sharpening fluxes. All regularization terms are constructed in divergence form, resulting in a method that conserves the mass of individual species as well as the mixture momentum and total energy.

Following Jain et al. [97], we consider the implied volume fraction transport equation for phase m, with the interface regularization volume fraction flux $(a_m)_k$,

$$\frac{\partial \phi_m}{\partial t} + u_k \frac{\partial \phi_m}{\partial x_k} = \frac{\partial (a_m)_k}{\partial x_k}. \tag{7.17}$$

In this work, this equation is not directly solved, because the volume fraction is closed during the pressure and temperature equilibration process (Section 7.2.4), but the action of this volume fraction flux is consistently incorporated into the system through the coupling terms with the other governing equations. We employ the coupling terms proposed by Jain et al. [97] for the mass, momentum, and energy equations, and propose new consistent coupling terms for the kinematic equations.

Using the relationship of density (ρ) and mass fraction (Y_m) to component density (ρ_m) and volume fraction (ϕ_m) for material m ($\rho Y_m = \rho_m \phi_m$, with no sum on repeated m), we can describe the interface regularization term for each material mass transport equation,

$$J_m = \frac{\partial (a_m)_k \, \rho_m}{\partial x_k}, \qquad \text{with no sum on repeated } m. \tag{7.18}$$

Consistent regularization terms for the momentum and energy equations follow,

$$F_i = \sum_m \left(\frac{\partial (a_m)_k \, \rho_m u_i}{\partial x_k} \right) \qquad \text{and} \qquad H = \sum_m \frac{\partial}{\partial x_k} \left\{ (a_m)_k \left[\frac{1}{2} \rho_m u_j u_j + (\rho h)_m \right] \right\}, \tag{7.19}$$

in which the enthalpy of species m is described,

$$(\rho h)_m = e_m \rho_m + p_m, \qquad \text{with no sum on repeated } m. \tag{7.20}$$

Consistent regularization terms for the kinematic equations take the form

$$K_{ij}^e = \frac{1}{3} \frac{1}{\rho} g_{ij}^e \sum_m J_m \tag{7.21}$$

This is derived by considering a transport equation for $\det \mathbf{g}^e$:

$$\frac{D}{Dt}(\det \mathbf{g}^e) = \frac{\partial \det \mathbf{g}^e}{\partial \mathbf{g}^e} : \frac{D}{Dt}\mathbf{g}^e \tag{7.22}$$

where $\frac{D}{Dt}$ denotes the material derivative, and ":" denotes the tensor inner product. Using identities from tensor calculus and the properties of the multiplicative decomposition, this can be simplified to

$$\frac{D\rho}{Dt} = \rho(\mathbf{g}^e)^{-T} : \frac{D}{Dt}\mathbf{g}^e \tag{7.23}$$

Plugging in the transport equations for $\mathbf{g}^e$ and ρ, converting to index notation, and ignoring the other artificial terms, this becomes

$$-\rho\frac{\partial u_k}{\partial x_k} + \sum_m J_m = \rho(g^e)_{ji}^{-1}\left(-g_{ik}^e\frac{\partial u_k}{\partial x_j} + K_{ij}^e\right) \tag{7.24}$$

The terms involving velocity cancel, leaving

$$\sum_m J_m = \rho\left((g^e)_{ji}^{-1}K_{ij}^e\right) \tag{7.25}$$

This relationship can be satisfied in many ways, but the form employed here is chosen for simplicity. No sharpening term is required in the equations for plastic deformation because in the multiplicative decomposition, volume change is entirely described by $\mathbf{g}^e$.

The volume fraction interface regularization flux for phase m is described by

$$(a_m)_k = \Gamma\left[\underbrace{\epsilon\frac{\partial \phi_m}{\partial x_k}}_{\substack{\text{interface} \\ \text{diffusion}}} - \underbrace{s_m\,(\hat{n}_m)_k}_{\substack{\text{interface} \\ \text{sharpening}}}\right]L_m + \underbrace{\Gamma^*\epsilon D_b\frac{\partial \phi_m}{\partial x_k}}_{\substack{\text{out-of-bounds} \\ \text{diffusion}}}, \qquad \text{with no sum on repeated } m, \tag{7.26}$$

with the interface sharpening term

$$s_m = \begin{cases} (\phi_m - \phi_m^\epsilon)\left(1 - \sum_{\substack{n=1 \\ n\neq m}}^{M}\phi_n^\epsilon - \phi_m\right), & \text{for } \phi_m^\epsilon \leq \phi_m \leq 1 - \sum_{\substack{n=1 \\ n\neq m}}^{M}\phi_n^\epsilon \\ 0, & \text{else} \end{cases}, \tag{7.27}$$

in which ϕ_m^ϵ denotes the minimum allowable volume fraction for phase m; this floor promotes physically realizable solutions to the pressure and temperature equilibria, which would otherwise not be well behaved if the mass or volume fraction exceeded the physically realizable bounds between

zero and one. We assume $\phi_m^\epsilon = 1 \times 10^{-6}$ unless stated otherwise. The optional mask term,

$$L_m = \begin{cases} 1, & \text{for } \phi_m^\epsilon \leq \phi_m \leq \sum_{\substack{n=1 \\ n \neq m}}^{M} \phi_n^\epsilon \\ 0, & \text{else} \end{cases}, \tag{7.28}$$

localizes the interface diffusion and interface sharpening terms to the interface region, restricting the application of the non-compactly discretized terms to the interface region. Unlike the gradient-form approach, this mask in the divergence-form approach is not necessary for stability, as demonstrated by Jain et al. [97]. The interface normal vector for phase m is given by

$$(\hat{n}_m)_k = \frac{\partial \phi_m}{\partial x_k} \bigg/ \left| \frac{\partial \phi_m}{\partial x_i} \right|, \qquad \left| \frac{\partial \phi_m}{\partial x_i} \right| = \sqrt{\frac{\partial \phi_m}{\partial x_i} \frac{\partial \phi_m}{\partial x_i}}, \qquad \text{with no sum on repeated } m. \tag{7.29}$$

The out-of-bounds diffusivity, described by

$$D_b = \max_m \overline{\left[1 - \phi_m / (\phi_m^\epsilon)^b \right]}_{\text{no sum in } m} (1 - L_m), \qquad b = \frac{1}{2}, \tag{7.30}$$

maintains $\phi_m \gtrsim \phi_m^\epsilon$. The overbar denotes the same filtering operation as applied to the LAD diffusivities. A user-specified length scale, $\epsilon \approx \Delta x$, typically on the order of the grid spacing, controls the equilibrium thickness of the diffuse interface. The velocity scale, $\Gamma \approx u_{max}$ controls the timescale over which the interface diffusion and interface sharpening terms drive the interface thickness to equilibrium. The velocity scale for the out-of-bounds volume fraction diffusivity is also specified by the user, with $\Gamma^* \gtrsim \Gamma$. Volume fraction compatibility is enforced by requiring that $\sum_{m=1}^{M} (a_m)_k = 0$.

7.2.7 Quasi-conservative gradient-form approach to interface regularization

As with the divergence-form approach, the interface regularization in this approach is achieved with the use of diffusion and sharpening terms that balance each other. Therefore, this method also results in constant interface thickness during the simulation. Shukla et al. [180] discuss disadvantages associated with the divergence-form approach due to the numerical differentiation of the interface normal vector. The numerical error of these terms can lead to interface distortion and grid imprinting due to the anisotropy of the derivative scheme. Ideally, we would like to have a regularization method that is conservative and that does not require any numerical differentiation of the interface normal vector. However, starting with the assumption of conservation, for nonzero regularization flux, we see that numerical differentiation of the interface normal vector can only be avoided in the limit that the divergence of the interface normal vector goes to zero. This limit corresponds to the limit of zero interface curvature, which cannot be avoided in multidimensional problems.

Therefore, this illustrates that a conservative method cannot be constructed for multidimensional applications without requiring differentiation of the interface normal vector; the non-conservative property (undesirable) of the gradient-form approach is a necessary consequence of the circumvention of interface-normal differentiation (desirable). This is demonstrated below, in which the phase subscript has been dropped.

$$
\frac{\partial}{\partial x_k} \left\{ \Gamma \left[\epsilon \frac{\partial \phi}{\partial x_k} + \phi \left(1 - \phi \right) \hat{n}_k \right] \right\}
$$

$$
= \hat{n}_k \frac{\partial}{\partial x_k} \left\{ \Gamma \epsilon \left| \frac{\partial \phi}{\partial x_j} \right| \right\} + \left\{ \Gamma \epsilon \left| \frac{\partial \phi}{\partial x_j} \right| \right\} \frac{\partial \hat{n}_k}{\partial x_k} + \frac{\partial \Gamma \phi \left(1 - \phi \right)}{\partial x_k} \hat{n}_k + \Gamma \phi \left(1 - \phi \right) \frac{\partial \hat{n}_k}{\partial x_k} \qquad (7.31)
$$

$$
\xrightarrow[\nabla \cdot \vec{n} \to 0]{} \hat{n}_k \frac{\partial}{\partial x_k} \left\{ \Gamma \left[\epsilon \left| \frac{\partial \phi}{\partial x_j} \right| + \phi \left(1 - \phi \right) \right] \right\},
$$

in which the final expression is obtained in the limit of $\nabla \cdot \vec{n} \to 0$.

Following Shukla et al. [180], we arrive at an implied volume fraction transport equation for phase m, with the interface regularization volume fraction term α_m,

$$
\frac{\partial \phi_m}{\partial t} + u_k \frac{\partial \phi_m}{\partial x_k} = (n_m)_k \frac{\partial \alpha_m}{\partial x_k}. \qquad (7.32)
$$

Unlike the divergence-form approach, the gradient-form approach requires no numerical differentiation of interface normal vectors, but it consequently results in conservation error. Like the divergence-form approach, this volume fraction transport equation is not directly solved, because the volume fraction is closed during the pressure and temperature equilibration process (Section 7.2.4), but the action of the volume fraction regularization term is consistently incorporated into the system of equations for mass, momentum, energy, and kinematic quantities through quasi-conservative coupling terms.

We employ an interface regularization term for each component mass transport equation consistent with the interface regularization volume fraction term,

$$
J_m = (n_m)_k \frac{\partial \alpha_m \rho_m}{\partial x_k}, \qquad \text{with no sum on repeated } m. \qquad (7.33)
$$

Because of the assumption of pressure and temperature equilibrium (volume fraction is a derived variable—not an independent state variable), it is important to form mass transport regularization terms consistently with the desired volume fraction regularization terms. In the method of Tiwari et al. [196], the terms do not need to be fully consistent (e.g., the component density is assumed to be slowly varying); the terms only need to produce similar interface profiles in the limit of $\Gamma \to \infty$ [180], because the volume fraction is an independent state variable. Following the assumption of Tiwari et al. [196] that the velocity, specific energy, and kinematic variables (but not the mixture density) vary slowly across the interface, the stability of the method is improved by further relaxing

conservation of the coupled equations. For example, the consistent regularization term for the momentum equation reduces to

$$F_i = \sum_m (n_m)_k \frac{\partial \alpha_m \rho_m u_i}{\partial x_k} \approx \sum_m (n_m)_k \frac{\partial \alpha_m \rho_m}{\partial x_k} u_i. \tag{7.34}$$

Similarly, the consistent regularization term for the energy equation reduces to

$$H = \sum_m (n_m)_k \frac{\partial}{\partial x_k} \left[\alpha_m \rho_m \left(\frac{1}{2} u_j u_j + h_m \right) \right] \approx \sum_m (n_m)_k \frac{\partial \alpha_m \rho_m}{\partial x_k} \left(\frac{1}{2} u_j u_j + h_m \right). \tag{7.35}$$

As with the divergence method, consistent regularization terms for the kinematic equations take the form

$$K_{ij}^e = \frac{1}{3} \frac{1}{\rho} g_{ij}^e \sum_m J_m. \tag{7.36}$$

No sharpening terms are required for the plastic deformation.

The volume fraction interface regularization flux for phase m is defined by

$$\alpha_m = \Gamma \left(\underbrace{\epsilon \left| \frac{\partial \phi_m}{\partial x_i} \right|}_{\substack{\text{interface} \\ \text{diffusion}}} - \underbrace{s_m}_{\substack{\text{interface} \\ \text{sharpening}}} \right) \mathscr{L}_m, \qquad \text{with no sum on repeated } m. \tag{7.37}$$

The volume fraction out-of-bounds diffusion term employed in the divergence-form approach (Eq. 7.26) is also active in the gradient-form approach. The gradient-form discretization of this term (including an equivalent volume fraction out-of-bounds term in Eq. 7.37) exhibits poor stability away from the interface, whereas the divergence-form approach does not. Following Shukla et al. [180] and Tiwari et al. [196], a necessary mask term blends the interface regularization terms to zero as the volume fraction approaches the specified minimum or maximum, thereby avoiding instability of the method away from the interface, where the calculation of the surface normal vector may behave spuriously and lead to compounding conservation error,

$$\mathscr{L}_m = \begin{cases} \tanh \left[\left(\frac{s_m}{\phi_m^{\mathscr{L}}} \right)^2 \right], & \text{for } \phi_m^\epsilon \leq \phi_m \leq 1 - \sum_{\substack{n=1 \\ n \neq m}}^{M} \phi_n^\epsilon \\ 0, & \text{else} \end{cases}, \tag{7.38}$$

in which $\phi_m^{\mathscr{L}} \approx 1 \times 10^{-2}$ is a user-specified value controlling the mask blending function. Other variables are the same as defined in the context of the divergence-form approach.

7.2.8   Numerical method

The equations are discretized on an Eulerian Cartesian grid. Time advancement is achieved using a five-stage, fourth-order, Runge-Kutta method, with an adaptive time step based on a Courant–Friedrichs–Lewy (CFL) condition. Other than the interface regularization terms for the divergence-form approach, all spatial derivatives are computed using a high-resolution, penta-diagonal, tenth-order, compact finite-difference scheme described by [116]. This scheme is applied in the domain interior and near the boundaries in the cases of symmetry, anti-symmetry, or periodic boundary conditions. Otherwise, boundary derivatives are reduced to a fourth-order, one-sided, compact difference scheme.

The interface sharpening and interface diffusion regularization terms in the divergence-form approach are discretized using node-centered derivatives, for which the fluxes to be differentiated are formed at the faces (staggered locations); linear terms (e.g., ϕ_i) are interpolated from the nodes to the faces, where the nonlinear terms are formed [e.g., $\phi_{i+1/2}\left(1 - \phi_{i+1/2}\right)$]. Here, we refer to the finite-difference grid points as nodes. All variables are stored at the nodes (collocated). If the nonlinear fluxes are not formed at the faces, poor stability is observed for node-centered finite-difference schemes of both compact and non-compact varieties due to the nonlinear interface sharpening term (see Appendix A). A second-order scheme is used for discretization of the interface regularization terms throughout this work, with an exception in Section 7.3.1 where both second-order and sixth-order (non-compact) discretization schemes are examined for these terms. The second-order scheme recovers the finite-volume approach employed by Jain et al. [97], whereas the higher-order scheme provides increased resolution and formal accuracy; however, discrete conservation is not guaranteed. The sixth order explicit scheme used to compute first derivatives from nodes to faces or vice versa is

$$f'_i = \frac{9}{384}\frac{f_{i+5/2} - f_{i-5/2}}{5h} - \frac{25}{128}\frac{f_{i+3/2} - f_{i-3/2}}{3h} + \frac{225}{192}\frac{f_{i+1/2} - f_{i-1/2}}{h} \qquad (7.39)$$

The sixth order interpolation scheme used for node to face or vice versa is

$$\hat{f}_i = \frac{3}{256}(f_{i+5/2} + f_{i-5/2}) - \frac{25}{256}(f_{i+3/2} + f_{i-3/2}) + \frac{75}{128}(f_{i+1/2} + f_{i-1/2}) \qquad (7.40)$$

The out-of-bounds diffusion is discretized using the tenth-order pentadiagonal scheme for all interface regularization approaches.

Following the approach used in previous works employing high-order finite difference schemes combined with localized artificial diffusivity for shock-capturing, a spatial dealiasing filter is applied after each stage of the Runge-Kutta algorithm to each of the conservative and kinematic variables to remove the top 10% of the grid-resolvable wavenumber content [41, 39, 73]. This mitigates against aliasing errors and numerical instability in the high-wavenumber range, which is not accurately resolved by the spatial derivative scheme. The filter is computed using a high-resolution, penta-diagonal, eighth-order, compact Padé filter, with cutoff parameters described by [73].

7.3   Results

In this section, we present the simulation results and evaluate the performance of the methods using classical two-dimensional test cases, such as: (a) advection of an air bubble in water, (b) shock interaction with a helium bubble in air, (c) shock interaction and the collapse of an air bubble in water, and (d) Richtmyer-Meshkov instability (RMI) of a copper-aluminium interface. The simulation test cases in the present study were carefully selected to assess: (1) the conservation property of the method; (2) the accuracy of the method in maintaining the interface shape; and (3) the ability of the method in maintaining constant interface thickness throughout the simulation.

Some of these test cases have been extensively studied in the past and have been used to evaluate the performance of various interface-capturing and interface-tracking methods. Many studies look at these test cases to evaluate the performance of the methods in the limit of very fine grid resolutions. For example, a typical value of the grid size is on the order of hundreds of mesh points across the diameter of a single bubble/droplet. However, for practical application of these methods in the large-scale simulations of engineering interest—where there are thousands of droplets, e.g., in an atomization process—it is rarely affordable to use such fine grids to resolve a single droplet/bubble. Therefore, in this study, we examine these methods in the opposite limit of relatively coarse grid resolution. This limit is more informative of the true performance of these methods for practical applications. However, the simulation results at higher resolutions are also presented in Appendix B, for the test case in Section 7.3.2, for the completeness of this study, and also to highlight the differences between the coarse and the refined grid limits. All three diffuse-interface capturing methods are implemented in the PadéOps solver [74] to facilitate fair comparison of the methods with the same underlying numerical methods, thereby eliminating any solver/implementation-related bias in the comparison.

The first test case (Section 7.3.1) is the advection of an air bubble in water. This test case is chosen to evaluate the ability of the interface-capturing method to maintain the interface shape for long-time numerical integration and to examine the robustness of the method for high-density-ratio interfaces. It is known that the error in evaluating the interface normal accumulates over time and results in artificial alignment of the interface along the grid [38, 196]. This behavior is examined for each of the three methods. The second test case (Section 7.3.2) is the shock interaction with a helium bubble in air. This test case is chosen to evaluate the ability of the methods to conserve mass, to maintain constant interface thickness throughout the simulation, and to examine the behavior of the under-resolved features captured by the methods. The third test case (Section 7.3.3) is the shock interaction with an air bubble in water. This test case is chosen to evaluate the robustness of the method to handle strong-shock/high-density-ratio interface interactions. The fourth test case (Section 7.3.4) is the RMI of a copper–aluminum interface. This test case is chosen to illustrate the applicability of the methods to simulate interfaces between solid materials with strength, to examine the conservation properties of the methods in the limit of high interfacial curvature, to examine the

ability of the methods to maintain constant interface thickness, and to assess the behavior of the under-resolved features captured by the methods.

For all the problems in this work, the mass fractions are initialized using the relations $Y_1 = \phi_1 \rho_1 / \rho$ and $Y_2 = 1 - Y_1$. To evaluate the mass-conservation property of a method, the total mass, m_k, of the phase k is calculated as

$$m_k = \int_\Omega \rho Y_k dv, \tag{7.41}$$

where the integral is computed over the computational domain Ω. To evaluate the ability of a method to maintain constant interface thickness, we define a new parameter—the interface-thickness indicator (l)—as

$$l = \left(\frac{1}{\hat{n} \cdot \vec{\nabla}\phi} \right), \tag{7.42}$$

and compute the maximum and average interface thicknesses in the domain, using

$$l_{max} = \max_{0.45 \leq \phi \leq 0.55} (l), \qquad l_{avg} = \langle l \rangle_{0.45 \leq \phi \leq 0.55}. \tag{7.43}$$

respectively, where $\langle \cdot \rangle$ denotes an averaging operation. The range for ϕ is used to ensure that the interface thickness is evaluated around the $\phi = 0.5$ isocontour because the quantity l is most accurate in this region and goes to ∞ as $\phi \to 0, 1$. Note that, occasionally, l can become very large, within the region $0.45 \leq \phi \leq 0.55$, when there is breakup due to the presence of a saddle point in the ϕ field at the location of rupture. These unphysical values of l show up in the computed l_{max} values and are removed during the post-processing step by plotting a moving average of 5 local minima of l_{max}. The unphysical values of l, on the other hand, have only a small effect on the computed l_{avg} values due to the averaging procedure.

7.3.1   Advection of an air bubble in water

This section examines advection of a circular air bubble in water. A one-dimensional version of this test case has been extensively studied before, and has been previously used as a test of robustness of the method using various diffuse-interface methods in Saurel and Abgrall [174], Allaire et al. [6], Murrone and Guillard [140], Johnsen and Ham [100], Saurel et al. [176], Johnsen and Colonius [99], Coralic and Colonius [42], Beig and Johnsen [25], Capuano et al. [36], and using a THINC method in Shyue and Xiao [182]. A two-dimensional advection of a bubble/drop has also been studied using a weighted-essentially non-oscillatory (WENO) and targeted-essentially non-oscillatory (TENO) schemes in Haimovich and Frankel [80].

In the current study, this test case is used to evaluate the ability of the methods in maintaining interface-shape for long-time integrations and as a test of robustness of the methods for high-density-ratio interfaces. Both phases are initialized with a uniform advection velocity. The problem domain spans $(0 \leq x \leq 1; 0 \leq y \leq 1)$, with periodic boundary conditions in both dimensions. The domain

is discretized on a uniform Cartesian grid of size $N_x = 100$ and $N_y = 100$. The bubble has a radius of $25/89$ and is initially placed at the center of the domain. The material properties for the water medium used in this test case are $\gamma_1 = 4.4$, $\rho_1 = 1.0$, $p_{\infty 1} = 6 \times 10^3$, $\mu_1 = 0$, and $\sigma_{Y1} = 0$. The material properties for the air medium used in this test case are $\gamma_2 = 1.4$, $\rho_2 = 1 \times 10^{-3}$, $p_{\infty 2} = 0$, $\mu_2 = 0$, and $\sigma_{Y2} = 0$, where $\gamma_k, \rho_k, p_{\infty k}, \mu_k$, and σ_{Yk} are the ratio of specific heats, density, stiffening pressure, shear modulus, and yield stress of phase k, respectively.

The initial conditions for the velocity, pressure, volume fraction, and density are

$$u = 5, \quad v = 0, \quad p = 1, \quad \phi_1 = \phi_1^{\epsilon} + (1 - 2\phi_1^{\epsilon}) f_\phi, \quad \phi_2 = 1 - \phi_1, \quad \rho = \phi_1 \rho_1 + \phi_2 \rho_2, \qquad (7.44)$$

respectively, in which the volume fraction function, f_ϕ, is given by

$$f_\phi = \frac{1}{2} \left\{ 1 - \mathrm{erf} \left[\frac{625/7921 - (x - 1/2)^2 - (y - 1/2)^2}{3 \Delta x} \right] \right\}. \qquad (7.45)$$

For this problem, the interface regularization length scale and the out-of-bounds velocity scale are defined by $\epsilon = \Delta x = 1.0 \times 10^{-2}$ and $\Gamma^* = 5.0$, respectively.

The simulation is integrated for a total physical time of $t = 1$ units, and the bubble at this final time is shown in Figure 7.1, facilitating comparison among the LAD, divergence-form, and gradient-form methods. All three methods perform well and are stable for this high-density-ratio case. The consistent regularization terms included in the momentum and energy equations are necessary to maintain stability. The divergence-form approach results in relatively faster shape distortion compared to the LAD and gradient-form approaches. This shape distortion is due to the accumulation of grid-induced anisotropic error resulting from numerical differentiation of the interface normal vector, which is required in the divergence-form approach but not the other approaches. A similar behavior of interface distortion was seen when the velocity was halved and the total time of integration was doubled, thereby confirming that this behavior is reproducible for a given flow-through time (results not shown).

Two possible ways to mitigate the interface distortion are by refining the grid or by using a higher-order scheme for the interface-regularization terms. Because we are interested in the limit of coarse grid resolution, we study the effect of using an explicit sixth-order finite-difference scheme to discretize the interface regularization terms. As described in Section 2.8, finite-difference schemes may be used to discretize the interface regularization terms—without resulting in spurious behavior—if the nonlinear interface sharpening and the counteracting diffusion terms are formed at the grid faces (staggered locations), from which the derivatives at the grid points (nodes) may be calculated. Comparing the second-order and sixth-order schemes for the interface regularization terms of the divergence-form approach, the final state of the advecting bubble is shown in Figure 7.2. The interface distortion is significantly reduced using the sixth-order scheme. More recently, Jain
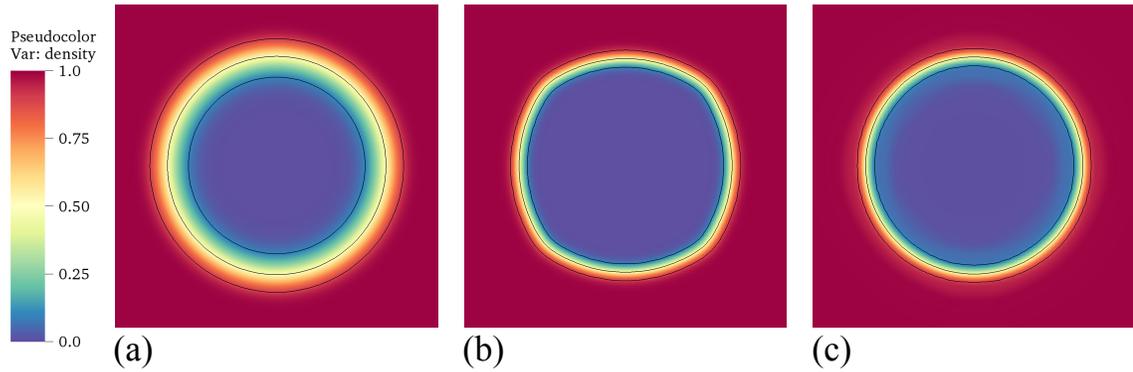
FIGURE 7.1: Comparison of the final state of the bubble after five flow-through times using (a) LAD approach, (b) divergence-form approach, and (c) gradient-form approach. The three solid black lines denote the isocontours of the volume fraction values of 0.1, 0.5, and 0.9, representing the interface region.
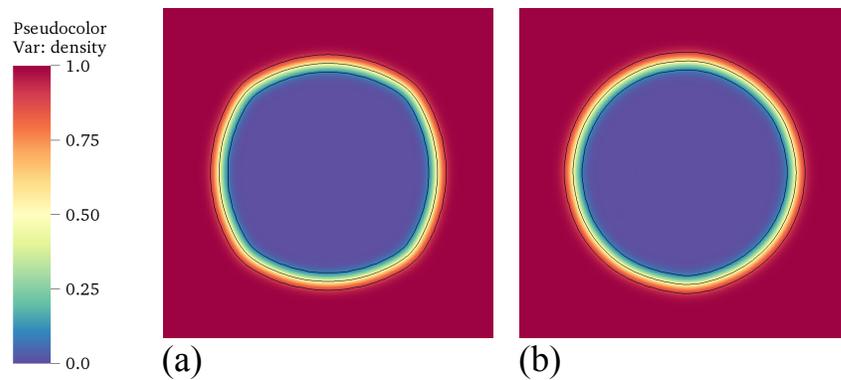
FIGURE 7.2: Comparison of the state of the bubble after five flow-through times using the divergence-form approach with (a) second-order scheme and (b) sixth-order scheme. The three solid black lines denote the isocontours of the volume fraction values of 0.1, 0.5, and 0.9, representing the interface region.
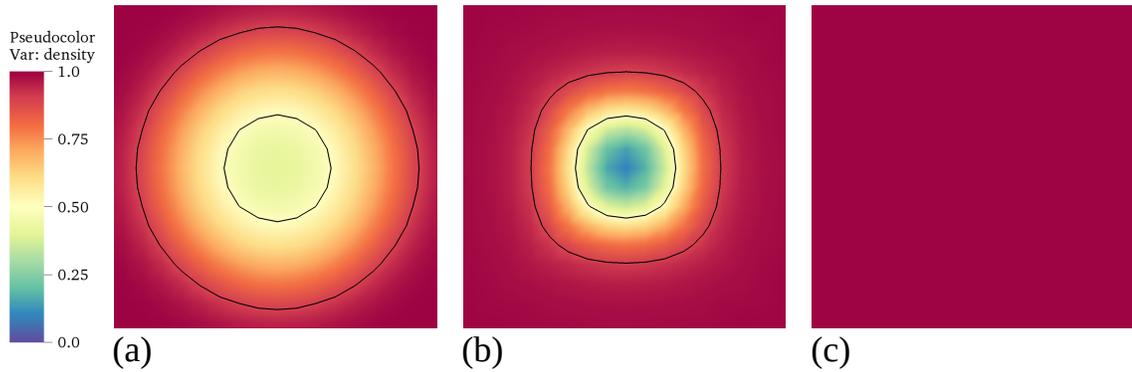
FIGURE 7.3: Comparison of the final state of the bubble on a coarse grid after five flow-through times using (a) LAD approach, (b) divergence-form approach, and (c) gradient-form approach. The two solid black lines denote the isocontours of the volume fraction values of 0.5 and 0.9, representing the interface region.

[92] proposed an alternative accurate conservative phase-field/diffuse-interface method, an ACDI method, that is not only conservative and maintains constant interface thickness throughout the simulation, but also results in higher accuracy by maintaining a much sharper interface and by eliminating the interface distortion.

Since the focus of the current work is on the evaluation of methods in a relatively coarser grid, we repeat the simulation of advection of an air bubble in water by scaling down the problem (in length and time) by a factor of 10 without changing the number of grid points. The domain length is kept the same, but the new bubble radius is 2.5/89, and the simulation is integrated for a total physical time of $t = 0.1$ units. At this resolution, the bubble has ≈ 5 grid points across its diameter, which represents a more realistic scenario that is encountered in large-scale engineering simulations.

The bubble at the final time is shown in Figure 7.3, for all the three methods. Similar to the more refined case above with the second-order finite-volume scheme, the divergence-form approach results in relatively faster shape distortion compared to the LAD method. Whereas, the gradient-form approach results in apparent complete loss of the bubble. Comparing this result with the refined simulation in Figure 7.1, this observation of mass loss on coarse grids is in good agreement with our hypothesis that the conservation error is proportional to the local interface curvature and the under-resolved features are more prone to being lost due to the non-conservative nature of the method. This makes the gradient-form approach unsuitable for large-scale engineering simulations where it is only possible to afford a couple of grids points across a bubble/drop. To quantify the amount of mass loss with the gradient-form approach, the bubble mass is plotted against time for all three methods in Figure 7.4. Interestingly, around 40% of the bubble mass is lost during the early time in the simulation and then the bubble mass saturates. This is due to the traces of mass of bubble that is still present in the domain, that is not under-resolved, after all the fine features are lost due to the conservation error.
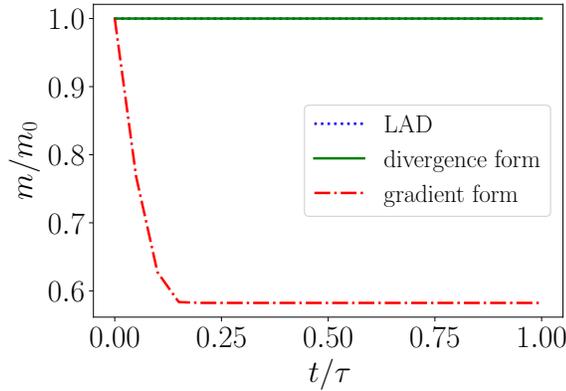
FIGURE 7.4: Plot of total mass, m, of the bubble by various methods, where m_0 is the mass at time $t = 0$.

7.3.2   Shock interaction with a helium bubble in air

This section examines the classic test case of a shock wave traveling through air followed by an interaction with a stationary helium bubble. This case has been extensively studied using various numerical methods and models, such as a front-tracking method in Terashima and Tryggvason [194]; an arbitrary-Lagrangian Eulerian (ALE) method in Daude et al. [47]; anti-diffusion interface-capturing method in So et al. [184]; a ghost-fluid method (GFM) in Fedkiw et al. [61], Bai and Deng [12]; a LAD diffuse-interface approach in Cook [40]; a gradient-form diffuse-interface approach in Shukla et al. [180], and other diffuse-interface methods that implicitly capture the interface (no explicit interface-capturing method) using a WENO scheme in Johnsen and Colonius [99], Coralic and Colonius [42], using a TENO scheme in Haimovich and Frankel [80], and using a WCNS scheme in Wong and Lele [210]. This test case has also been simulated with an adaptive-mesh-refinement technique in Quirk and Karni [167] where a refined grid is used around the interface to improve the accuracy. More recently, this test case has also been studied in a three-dimensional setting in Deng et al. [50].

To examine the interface regularization methods, we model this problem without physical species diffusion; therefore, the interface regularization methods for immiscible phases are applicable, because no physical molecular mixing should be exhibited by the underlying numerical model. The use of immiscible interface-capturing methods to model the interface between the gases in this problem is also motivated by the experiments of Haas and Sturtevant [79]. In these experiments, the authors use a thin plastic membrane to prevent molecular mixing of helium and air.

The problem domain spans $(-2 \leq x \leq 4; 0 \leq y \leq 1)$, with periodic boundary conditions in the y direction. A symmetry boundary is applied at $x = 4$, representing a perfectly reflecting wall, and a sponge boundary condition is applied over $(-2 \leq x \leq -1.5)$, modeling a non-reflecting free boundary. The problem is discretized on a uniform Cartesian grid of size $N_x = 600$ and $N_y = 100$.

The bubble has a radius of $25/89$ and is initially placed at the location $(x = 0, y = 1/2)$. The material properties for the air medium are described by $\gamma_1 = 1.4$, $\rho_1 = 1.0$, $p_{\infty 1} = 0$, $\mu_1 = 0$, and $\sigma_{Y1} = 0$. The material properties for the helium medium are described by $\gamma_2 = 1.67$, $\rho_2 = 0.138$, $p_{\infty 2} = 0$, $\mu_2 = 0$, and $\sigma_{Y2} = 0$.

The initial conditions for the velocity, pressure, volume fraction, and density are

$$u = u^{(2)} f_s + u^{(1)} \left(1 - f_s\right), \quad v = 0, \quad p = p^{(2)} f_s + p^{(1)} \left(1 - f_s\right),$$
$$\phi_1 = \phi_1^\epsilon + \left(1 - 2\phi_1^\epsilon\right) f_\phi, \quad \phi_2 = 1 - \phi_1, \quad \rho = \left(\phi_1 \rho_1 + \phi_2 \rho_2\right) \left[\rho^{(2)}/\rho^{(1)} f_s + \left(1 - f_s\right)\right],$$
(7.46)

respectively, in which the volume fraction function, f_ϕ, and the shock function, f_s, are given by

$$f_\phi = \frac{1}{2} \left\{1 - \mathrm{erf}\left[\frac{625/7921 - x^2 - (y - 1/2)^2}{\Delta x}\right]\right\} \quad \text{and} \quad f_s = \frac{1}{2} \left[1 - \mathrm{erf}\left(\frac{x + 1}{2\Delta x}\right)\right], \quad (7.47)$$

respectively, with jump conditions across the shock for velocity $\left(u^{(1)} = 0; u^{(2)} = 0.39473\right)$, density $\left(\rho^{(1)} = 1, \rho^{(2)} = 1.3764\right)$, and pressure $\left(p^{(1)} = 1; p^{(2)} = 1.5698\right)$. For this problem, the interface regularization length scale and the out-of-bounds velocity scale are defined by $\epsilon = \Delta x = 0.01$ and $\Gamma^* = 2.5$, respectively.

The interaction of the shock with the helium bubble and the eventual breakup of the bubble are shown in Figure 7.5, with depictions of the evolution at various times, for LAD, divergence-form, and gradient-form approaches. The bubble can be seen to undergo breakup at an approximate (non-dimensional) time of $t = 2.5$. After this time, the simulation cannot be considered physical because of the under-resolved processes associated with the breakup and the lack of explicit subgrid models for these processes; each interface regularization approach treats the under-resolved processes differently. Therefore, there is no consensus on the final shape of the bubble among the three methods. Yet, a qualitative comparison between the three methods can still be made using the results presented in Figure 7.5.

Using the LAD approach, the interface diffuses excessively in the regions of high shear, unlike the divergence-form and gradient-form approaches, where the interface thickness is constant throughout the simulation. However, using the LAD approach, the interface remains sharp in the regions where there is no shearing. To quantify the amount of interface diffusion, the interface-thickness indicator [l of Eq. (7.43)] is plotted in Figure 7.6 for the three methods. The average thickness, l_{avg}, increases slightly for the LAD method around $t/\tau \approx 2$, but the maximum interface thickness, l_{max}, increases almost 15 times for the LAD method, whereas it remains on the order of one for the other two methods. This demonstrates a deficiency of the LAD approach for problems that involve significant shearing at an interface that is not subjected to compression.

Furthermore, the behavior of bubble breakup is significantly different among the various methods. Depending on the application, any one of these methods may or may not result in an appropriate

FIGURE 7.5: Comparison of the bubble shapes at different times for the case of the shock/helium-bubble-in-air interaction using various interface-capturing methods. The three solid black lines denote the isocontours of the volume fraction values of 0.1, 0.5, and 0.9, representing the interface region.
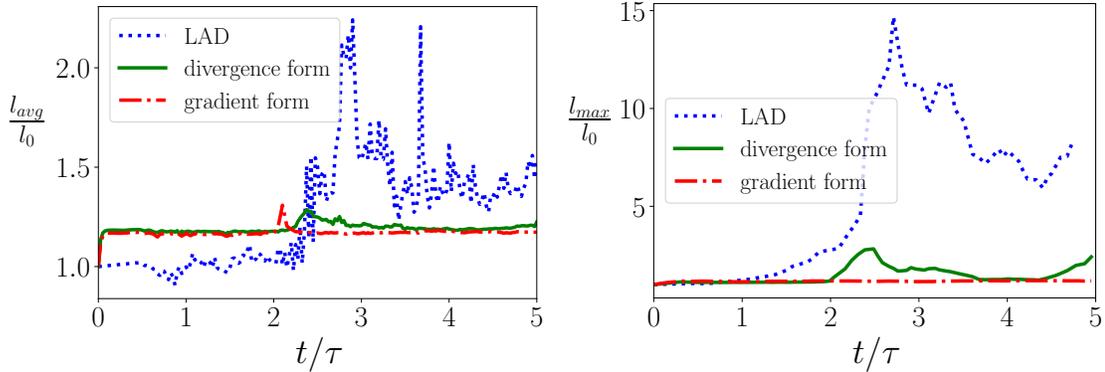
FIGURE 7.6: Comparison of the interface-thickness indicator, l, by various methods, where l_0 is the maximum interface thickness at time $t = 0$. (a) Average interface thickness l_{avg}. (b) Maximum interface thickness l_{max}. To exclude unphysical spikes in l_{max} during breakup events, a moving average of 5 local minima of l_{max} is plotted.

representation of the under-resolved processes. However, for the current study that involves modeling interfaces between immiscible fluids, the grid-induced breakup of the divergence-form approach may be more suitable than the diffusion of the fine structures in the LAD approach or the premature loss of fine structures and associated conservation error of the gradient-form approach. For the LAD approach, the thin film formed at around time $t = 2.1$ does not break; rather, it evolves into a thin region of well-mixed fluid. This behavior may be considered unphysical for two immiscible fluids, for which the physical interface is infinitely sharp in a continuum sense; this behavior would be more appropriate for miscible fluids. For the divergence-form approach, the thin film forms satellite bubbles, which is expected when there is a breakage of a thin ligament between droplets or bubbles due to surface-tension effects. However, this breakup may not be considered completely physical without any surface-tension forces, because the breakup is triggered by the lack of grid support. For the gradient-form approach, the thin film formed at around time $t = 2.1$ breaks prematurely and disappears with no formation of satellite bubbles, and the mass of the film is lost to conservation error.

In Figure 2 of Shukla et al. [180], without the use of interface regularization terms, the interface thickness is seen to increase significantly. Their approach without interface regularization terms is most similar to our LAD approach, because the LAD approach does not include any sharpening terms. Therefore, comparing these results suggests that the thickening of the interface in their case was due to the use of the more dissipative Riemann-solver/reconstruction scheme. The results from the gradient-form approach also match well with the results of the similar method shown in Figure 2 of Shukla et al. [180], which further verifies our implementation. Finally, there is no consensus on the final shape of the bubble among the three methods, which is to be expected, because there are no surface-tension forces and the breakup is triggered by the lack of grid resolution.
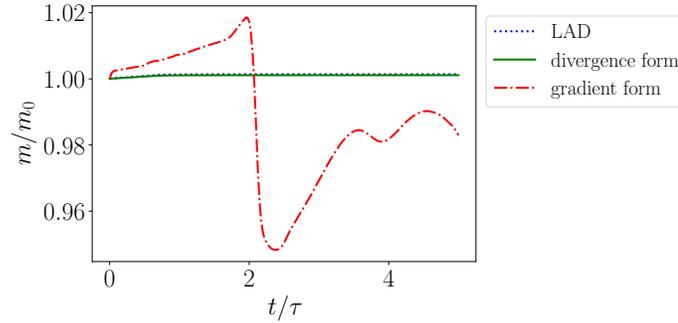
FIGURE 7.7: Plot of total mass, m, of the helium bubble by various methods, where m_0 is the mass at time $t = 0$.

To further quantify the amount of mass lost or gained, the total mass of the bubble is computed using Eq. (7.41) and is plotted over time in Figure 7.7. The mass of the bubble is conserved for the LAD and divergence-form approaches, but is not conserved for the gradient-form approach, as expected. The loss of mass is observed to be largest when the bubble is about to break, for the gradient form approach. This is because the mass-conservation error in the gradient-form approach is proportional to the local curvature, as described in Section 2.7. Therefore, at the onset of breakup, thin film rupture is different from the other two methods, and the satellite bubbles are absent.

7.3.3 Shock interaction with an air bubble in water

This section examines a shock wave traveling through water followed by an interaction with a stationary air bubble. The material properties are the same as those described in Section 7.3.1. This test case is based on the experiments in Bourne and Field [29] and has been widely used as a validation case for various numerical methods and models such as a front-tracking method in Terashima and Tryggvason [194]; a level-set method in Hu and Khoo [89], Nourgaliev et al. [146]; a ghost-fluid method in Bai and Deng [12]; a volume-of-fluid method in Bo and Grove [27]; an implicit diffuse-interface method with a Godunov scheme in Ansari and Daramizadeh [8], and with a TENO scheme in Haimovich and Frankel [80]; and with a gradient-form diffuse-interface approach in Shukla et al. [180], Shukla [179].

The initial conditions for the velocity, pressure, volume fraction, and density are as given in Eq. 7.46, in which the volume fraction function, f_ϕ, and the shock function, f_s, are given by,

$$f_\phi = \frac{1}{2}\left\{1 - \text{erf}\left[\frac{1 - (x - 2.375)^2 - (y - 2.5)^2}{\Delta x}\right]\right\} \quad \text{and} \quad f_s = \frac{1}{2}\left[1 - \text{erf}\left(\frac{x + 1}{10\Delta x}\right)\right], \quad (7.48)$$

respectively, with jump conditions across the shock for velocity $\left(u^{(1)} = 0;\ u^{(2)} = 68.5176\right)$, density $\left(\rho^{(1)} = 1,\ \rho^{(2)} = 1.32479\right)$, and pressure $\left(p^{(1)} = 1;\ p^{(2)} = 19150\right)$. The problem domain spans

$(-2 \leq x \leq 8; \, 0 \leq y \leq 5)$, with periodic boundary conditions in the y direction. A symmetry boundary is applied at $x = 8$, representing a perfectly reflecting wall, and a sponge boundary condition is applied over $(-2 \leq x \leq -1.5)$, modeling a non-reflecting free boundary. The problem is discretized on a uniform Cartesian grid of size $N_x = 400$ and $N_y = 200$.

For this problem, the artificial bulk viscosity, artificial thermal conductivity, artificial diffusivity, interface regularization length scale, interface regularization velocity scale, and out-of-bounds velocity scale are defined by $C_\beta = 20$, $C_\kappa = 0.1$, $C_D = 20$, $\epsilon = \Delta x = 2.5 \times 10^{-2}$, $\Gamma = 2.0$, and $\Gamma^* = 0.0$, respectively. A fourth-order, penta-diagonal, Padé filter is employed for dealiasing in this problem to improve the stability of the shock/bubble interaction. The linear system defining this filter is given by

$$\hat{f}_i + \alpha(\hat{f}_{i+1} + \hat{f}_{i-1}) + \frac{1 - 2\alpha}{14}(\hat{f}_{i+2} + \hat{f}_{i-2}) = \frac{4 + 6\alpha}{7}f_i + \frac{2 + 3\alpha}{7}(f_{i+1} + f_{i-1}) \qquad (7.49)$$

where $\alpha = 0.499$.

Notably, for this problem, the LAD in the mass equations is also necessarily included in the divergence-form and gradient-form approaches to maintain stability. The latter approaches become unstable for this problem for large Γ (the velocity scale for interface regularization). Figure 7.8 describes the evolution in time of the shock/bubble interaction and the subsequent bubble collapse. There is no significant difference between the various regularization methods for this problem. The similarity is due to the short convective timescale of the flow relative to the maximum stable timescale of the volume fraction regularization methods; effectively, all methods remain qualitatively similar to the LAD approach.

7.3.4   Richtmyer–Meshkov instability of a copper–aluminum interface

This section examines a shock wave traveling through copper followed by an interaction with a sinusoidally distorted copper–aluminum material interface. Though this problem has not been as widely studied as the previous examples, it is included to demonstrate how interface regularization methods perform when extended to problems involving elastic-plastic deformation at material interfaces. Such deformation may arise in impact welding, where interfacial instabilities are known to develop as metal plates impact and shear [142]; as well as material characterization at high strain rates, which typically employ a metal-gas configuration of the Richtmyer-Meshkov instability [51]. The copper-aluminum variant of this problem was previously studied by [120], who used a level-set method combined with the modified ghost-fluid method to set boundary conditions at material interfaces. This problem was also studied by [188] and [5], and the results presented here are an extension of that work.

The problem domain spans $(-2 \leq x \leq 4; \, 0 \leq y \leq 1)$, with periodic boundary conditions in the y direction. A symmetry boundary is applied at $x = 4$, representing a perfectly reflecting wall,
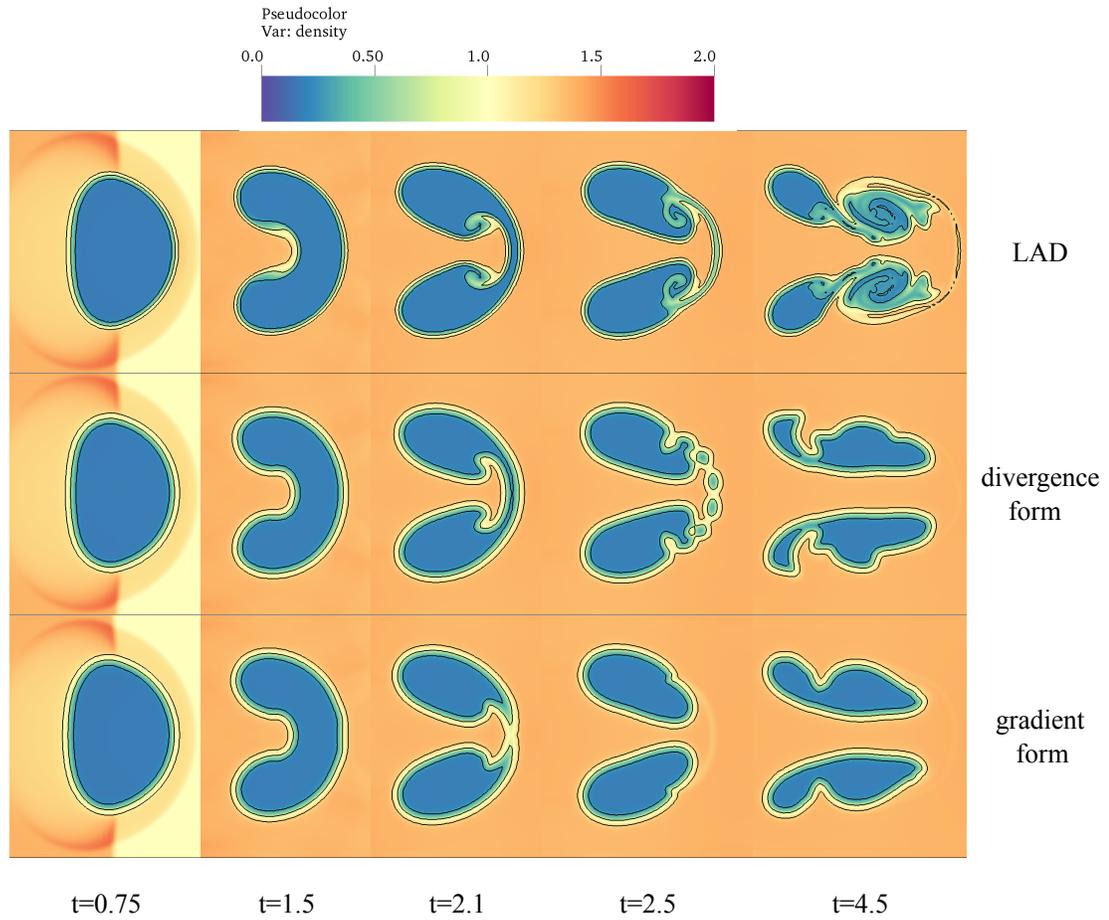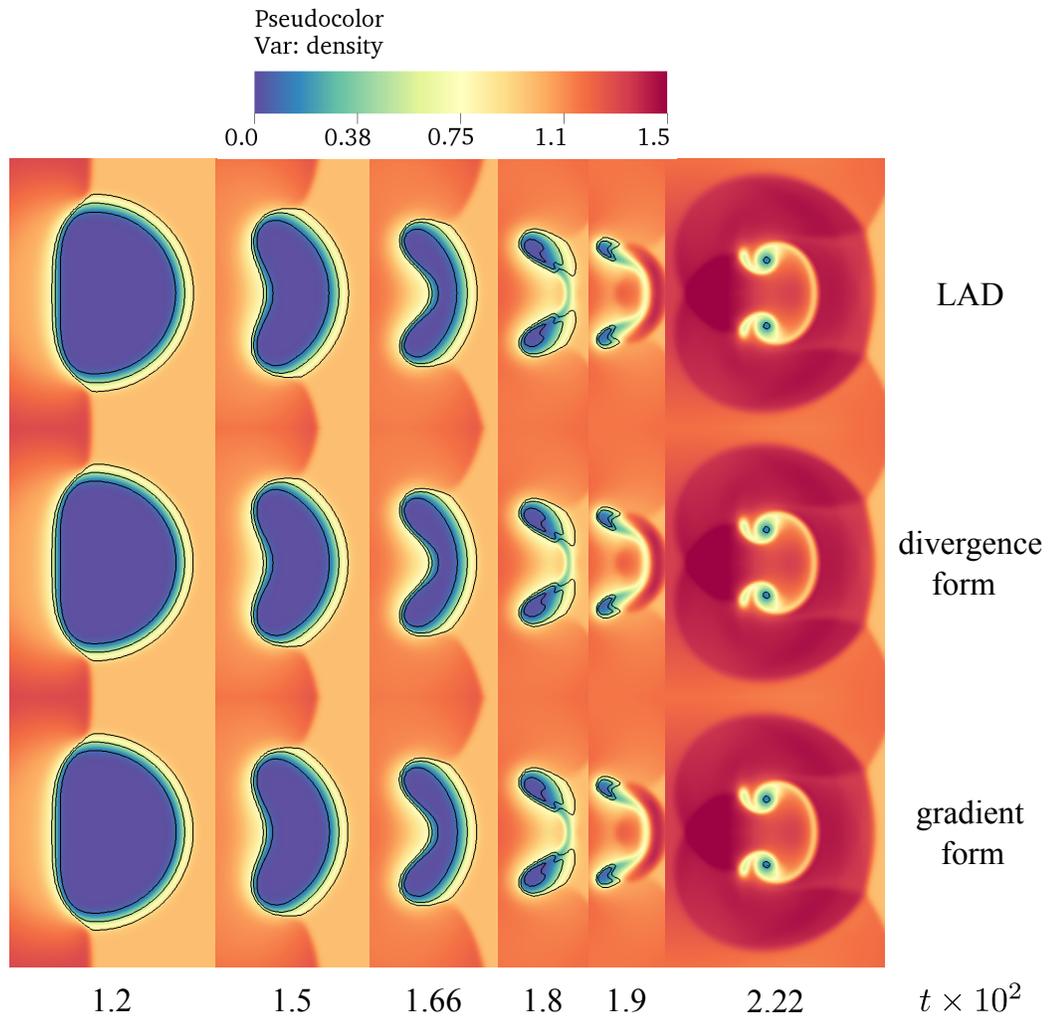
FIGURE 7.8: Comparison of the bubble shapes at different times for the case of shock/air-bubble-in-water interaction using various interface-capturing methods. The three solid black lines denote the isocontours of the volume fraction values of 0.1, 0.5, and 0.9, representing the interface region.
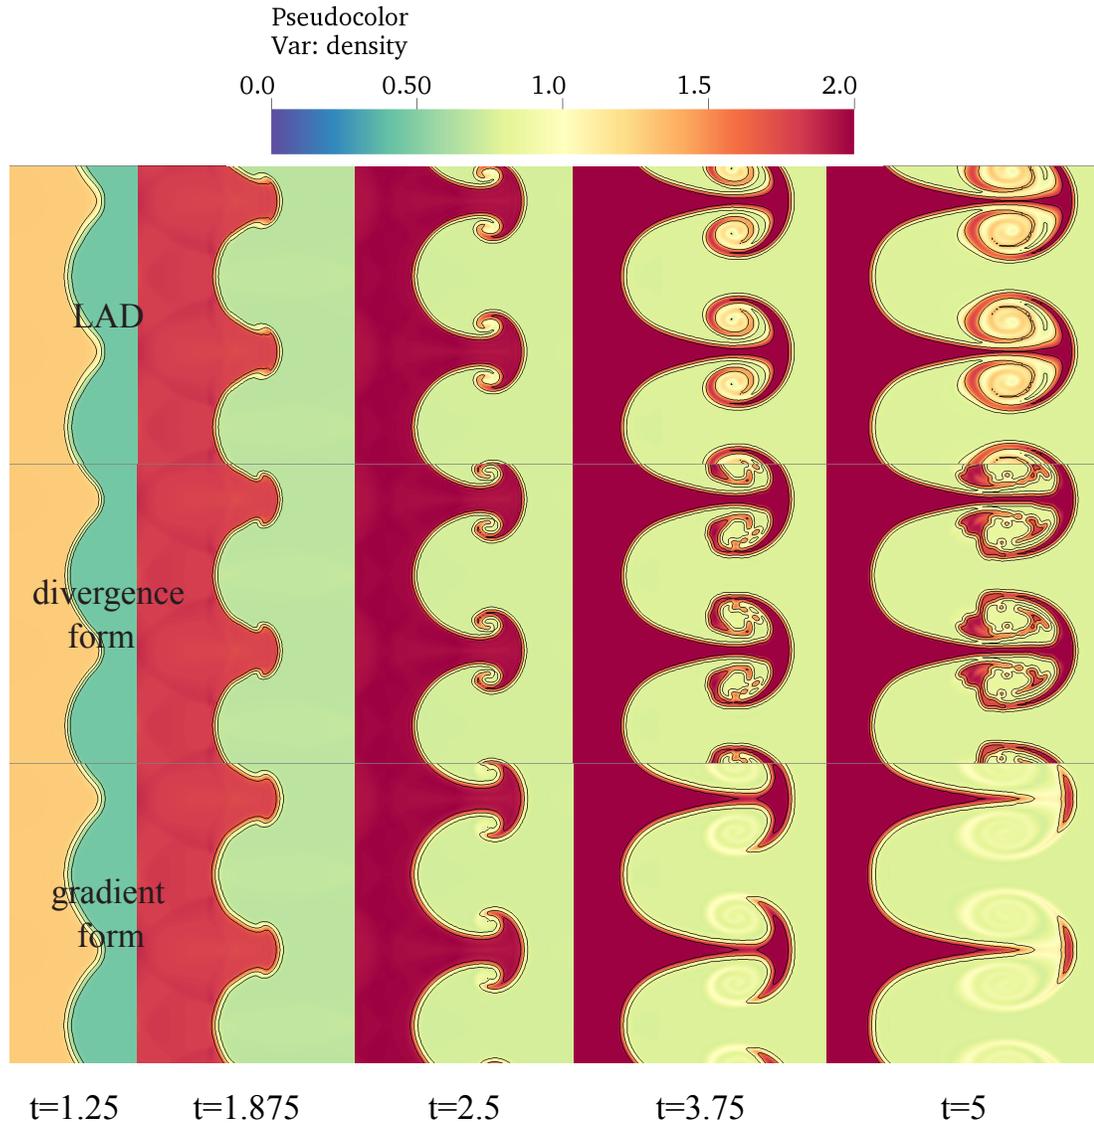
FIGURE 7.9: Comparison of the copper–aluminum interface shapes at different times for the Cu-Al RMI case using various interface-capturing methods. The three solid black lines denote the isocontours of the volume fraction values of 0.1, 0.5, and 0.9, representing the interface region.
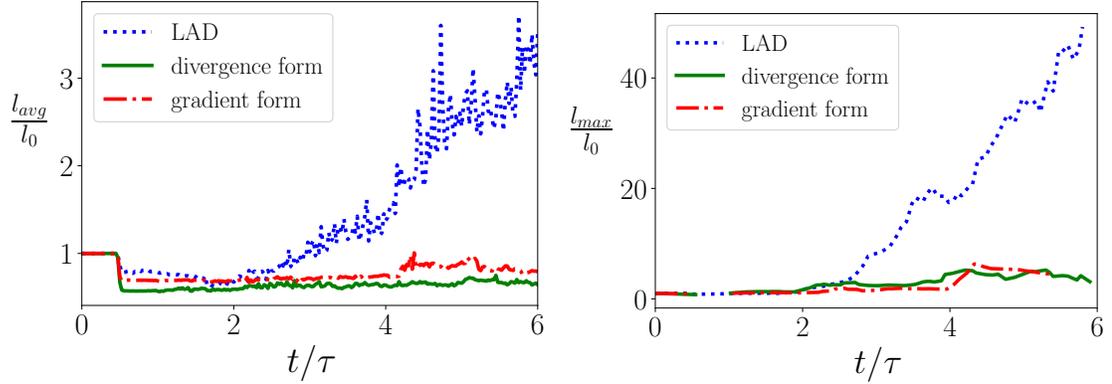
FIGURE 7.10: Comparison of the interface-thickness indicator, l, by various methods, where l_0 is the maximum interface thickness at time $t = 0$. (a) Average interface thickness l_{avg}. (b) Maximum interface thickness l_{max}. To exclude unphysical spikes in l_{max} during breakup events, a moving average of 5 local minima of l_{max} is plotted.

and a sponge boundary condition is applied over $(-2 \leq x \leq -1.5)$, modeling a non-reflecting free boundary. The problem is discretized on a uniform Cartesian grid of size $N_x = 768$ and $N_y = 128$. The material properties for the copper medium are described by $\gamma_1 = 2.0$, $\rho_1 = 1.0$, $p_{\infty 1} = 1.0$, $\mu_1 = 0.2886$, and $\sigma_{Y1} = 8.79 \times 10^{-4}$. The material properties for the aluminum medium are described by $\gamma_2 = 2.088$, $\rho_2 = 0.3037$, $p_{\infty 2} = 0.5047$, $\mu_2 = 0.1985$, and $\sigma_{Y2} = 2.176 \times 10^{-3}$.

The initial conditions for the velocity, pressure, volume fraction, and density are

$$u = u^{(2)} f_s + u^{(1)} \left(1 - f_s\right), \quad v = 0, \quad p = p^{(2)} f_s + p^{(1)} \left(1 - f_s\right),$$

$$\phi_1 = \phi_1^\epsilon + (1 - 2\phi_1^\epsilon) f_\phi, \quad \phi_2 = 1 - \phi_1, \quad \rho = (\phi_1 \rho_1 + \phi_2 \rho_2) \left[\rho^{(2)}/\rho^{(1)} f_s + (1 - f_s)\right], \tag{7.50}$$

respectively, in which the volume fraction function, f_ϕ, and the shock function, f_s, are given by

$$f_\phi = \frac{1}{2}\left(1 - \mathrm{erf}\left\{\frac{x - [2 + 0.4/(4\pi y)\sin(4\pi y)]}{3\Delta x}\right\}\right) \quad \text{and} \quad f_s = \frac{1}{2}\left[1 - \mathrm{erf}\left(\frac{x - 1}{2\Delta x}\right)\right], \tag{7.51}$$

respectively, with jump conditions across the shock for velocity $\left(u^{(1)} = 0; \, u^{(2)} = 0.68068\right)$, density $\left(\rho^{(1)} = 1, \, \rho^{(2)} = 1.4365\right)$, and pressure $\left(p^{(1)} = 5 \times 10^{-2}; \, p^{(2)} = 1.25\right)$. The kinematic tensors are initialized in a pre-strained state consistent with the material compression associated with shock initialization, assuming no plastic deformation has yet occurred, with

$$g_{ij} = g_{ij}^e = \begin{cases} \left[\rho^{(2)} f_s + \rho^{(1)} \left(1 - f_s\right)\right]/\rho_1, & \text{for } i = j = 1 \\ \delta_{ij}, & \text{else} \end{cases} \quad \text{and} \quad g_{ij}^p = \delta_{ij}. \tag{7.52}$$

For this problem, the interface regularization length scale and the out-of-bounds velocity scale for
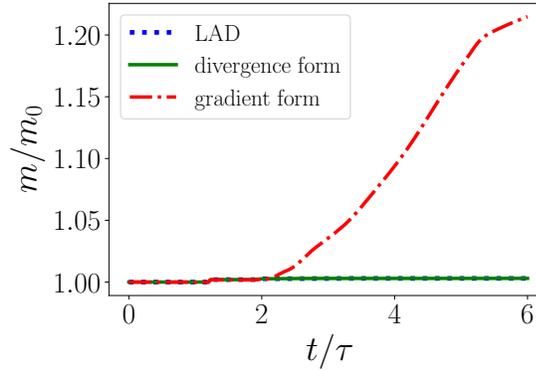
FIGURE 7.11: Plot of total mass, m, of aluminum by various methods, where m_0 is the mass at time $t = 0$.

the divergence form method are defined by $\epsilon = \Delta x/2 = 3.90625 \times 10^{-3}$ and $\Gamma^* = 1.0$, respectively. For the gradient form method, it is necessary for stability to use $\epsilon = 3\Delta x/4 = 5.859375 \times 10^{-3}$, $\Gamma^* = 1.0$, and $\phi_{min} = 1 \times 10^{-5}$.

The time evolution of the growth of the interface instability is shown in Figure 7.9. The simulation is integrated well into the nonlinear regime where the bubble (lighter medium) and the spike (heavier medium) have interpenetrated, forming mushroom-shaped structures with fine ligaments. The qualitative comparison between the methods in this test case is similar to that of the shock-helium-bubble interaction in air. With the LAD approach, the interface thickness increases with time, especially in the regions of high shear at the later stages. However, with the divergence-form and gradient-form approaches, the interface thickness is constant throughout the simulation. This is quantified by plotting the interface-thickness indicator [l of Eq. (7.43)] for each of the three methods in Figure 7.10. The average thickness, shown in Figure 7.10(a) shows a sharp drop in thickness at $t/\tau \approx 0.5$ when the shock passes through the interface. After this, the thickness remains small for both the gradient and divergence form methods, whereas with LAD the interface thickness grows gradually after $t/\tau \approx 2$, when the interface begins to roll up. Figure 7.10(b) shows the maximum interface thickness, which increases almost 60 times for the LAD method, whereas it stays on the order of one for the other two methods. This illustrates that the LAD method incurs significant artificial diffusion when the interface deformation cannot be resolved by the grid.

It is also evident from Figure 7.9 that the gradient-form approach results in significant copper mass loss, and the dominant mushroom structure formed in the nonlinear regime is completely lost. To quantify the amount of mass lost or gained, the total mass of the aluminum material [Eq. (7.41)] is plotted against time in Figure 7.11. The gradient-form approach results in significant gain in the mass of the aluminum material, up to 20%, as the grid can no longer resolve the increased interface curvature during roll-up. This makes it practically unsuitable for accurate interface representation for long-time numerical simulations. With the divergence-form approach, the breakage of the

ligaments to form metallic droplets can be seen in Figure 7.9.

7.4   Summary and concluding remarks

This work examines three diffuse-interface-capturing methods and evaluates their performance for the simulation of immiscible compressible multiphase fluid flows and elastic-plastic deformation in solids. The first approach is the *localized-artificial-diffusivity* method of Cook [39], Subramaniam et al. [188], and Adler and Lele [5], in which artificial diffusion terms are added to the individual phase mass fraction transport equations and are coupled with the other conservation equations. The second approach is the *gradient-form approach* that is based on the quasi-conservative method of Shukla et al. [180]. In this method, the diffusion and sharpening terms (together called regularization terms) are added to the individual phase volume fraction transport equations and are coupled with the other conservation equations [196]. The third approach is the *divergence-form approach* that is based on the fully conservative method of Jain et al. [97]. In this method, the diffusion and sharpening terms are added to the individual phase volume fraction transport equations and are coupled with the other conservation equations. In the present study, all of these interface regularization methods are used in conjunction with a four-equation, multicomponent mixture model, in which pressure and temperature equilibria are assumed among the various phases. The latter two interface regularization methods are commonly used in the context of a five-equation model, in which temperature equilibrium is not assumed.

The primary objective of this work is to compare these three methods in terms of their ability to: maintain constant interface thickness throughout the simulation; conserve mass of each of the phases, mixture momentum, and total energy; and maintain accurate interface shape for long-time integration. The secondary objective of this work is to extend these methods for modeling the interface between deforming solid materials with strength. The LAD method has previously been used for simulating material interfaces between solids with strength [188, 5]. Here, we introduce consistent corrections in the kinematic equations for the divergence-form and the gradient-form approaches to extend these methods for the simulation of interfaces between solids with strength.

We employ several test cases to evaluate the performance of the methods, including (1) advection of an air bubble in water, (2) shock interaction with a helium bubble in air, (3) shock interaction and the collapse of an air bubble in water, and (4) Richtmyer–Meshkov instability of a copper–aluminum interface. For the application of these methods to large-scale simulations of engineering interest, it is rarely practical to use hundreds of grid points to resolve the diameter of a bubble/drop. Therefore, we choose to study the limit of relatively coarse grid resolution, which is more representative of the true performance of these methods (results for the higher resolution limit can be found in Appendix B).

---

[1]The new ACDI method in Jain [92] overcomes this issue

Method	Conservation	Sharp interface	Shape preservation	Behavior of under-resolved ligaments and breakup
LAD	Yes	<u>No</u> (interface diffuses in the regions of high shear)	Yes	<u>Artificial diffusion</u> (fine-scale features artificially diffuse as they approach unresolved scales)
Divergence form	Yes	Yes	<u>No</u>[1] (interface aligns with the grid)	<u>Artificial breakup</u> (fine-scale features artificially break up as they approach unresolved scales)
Gradient form	<u>No</u> (under-resolved features will be lost)	Yes	Yes	<u>Artificial loss of mass</u> (fine-scale features are lost, due to conservation error, as they approach unresolved scales)

TABLE 7.1: Summary of the advantages and disadvantages of the three diffuse-interface capturing methods considered in this study: LAD method based on Cook [39], Subramaniam et al. [188], and Adler and Lele [5]; divergence-form approach based on Jain et al. [97]; and the gradient-form approach based on Shukla et al. [180] and Tiwari et al. [196]. The relative disadvantages of each approach and the different behaviors of under-resolved processes are underlined.

The performance of the three methods is summarized in Table 7.1. The LAD and the divergence-form approaches conserve mass, momentum, and energy, whereas the gradient-form approach does not. The mass-conservation error increases proportionately with the local interface curvature; therefore, fine interfacial structures will be lost during the simulation. The divergence-form and the gradient-form approaches maintain a constant interface thickness throughout the simulation, whereas the interface thickness of the LAD method increases in the regions of high shear due to the lack of interface sharpening terms to counter the artificial diffusion. The LAD and the gradient-form approaches maintain the interface shape for a long time compared to the divergence-form approach; however, the interface distortion of the divergence-form approach can be mitigated with the use of appropriately crafted higher-order schemes for the interface regularization terms. A new diffuse-interface method, that is also in divergence form, can be found in Jain [92] that is not only conservative and maintains a constant interface thickness throughout the simulation, but is also more accurate because it maintains a much sharper interface and eliminates the interface distortion.

For each method, the behavior of under-resolved ligaments and breakup features is unique. For the LAD approach, thin ligaments that form at the onset of bubble breakup (or in late-stage RMI) diffuse instead of rupturing. For the gradient-form approach, the ligament formation is not captured because of mass-conservation issues, which result in premature loss of these fine-scale features. For the divergence-form approach, the ligaments rupture due to the lack of grid support, acting like an artificial surface tension force that becomes significant at the grid scale.

For broader applications, the optimal method depends on the objectives of the study. These applications include (1) well-resolved problems, in which differences in the behavior of under-resolved features is not of concern, (2) applications involving interfaces between miscible phases, and (3) applications involving more complex physics, including regimes in which surface tension or molecular diffusion must be explicitly modeled and problems in which phase changes occur. We intend this demonstration of the advantages, disadvantages, and behavior of under-resolved phenomena exhibited by the various methods to be helpful, albeit being unphysical, in the selection of an interface-regularization method. These results also provide motivation for the development of subgrid models for multiphase flows.

7.5 Appendix A: Finite-difference operators for the divergence-form approach

The test case of shock interaction with a helium bubble in air is repeated for the divergence-form approach with the same parameters listed in Section 7.3.2. Here, the difference is in the numerical representation of the nonlinear interface-regularization terms. In Section 7.3.2, the interface-regularization fluxes are formed at the faces, as described in Section 7.2.8, which is consistent with the finite-volume implementation in Jain et al. [97]. Whereas, here, a second-order standard central finite-difference scheme is used instead.

The shock interaction with the helium bubble in air and the subsequent evolution of the bubble shape are shown in Figure 7.12. An unphysical wrinkling of the interface can be seen at the later stages of the bubble deformation. This behavior is consistent with the observations made by Shukla et al. [180], which motivated them to develop the gradient-form approach. However, discretizing the fluxes at the faces, Jain et al. [97] showed that this results in discrete balance between the diffusion and sharpening fluxes, thereby eliminating the spurious wrinkling of the interface as can be seen in Figure 7.5. In this work, this face-evaluated flux formulation has been succesfully extended for higher-order schemes and is presented in Section 7.2.8.

7.6 Appendix B: Grid refinement study

Throughout this work, the simulations are presented at coarse grid resolutions and the methods are evaluated in this coarse grid limit because this limit is more informative of the true performance of the methods. However, in this section, a grid refinement study is performed to evaluate and compare the methods in the refined grid limit for the completeness of this study, and to highlight the differences between the coarse and refined grid limits.

The test case of a shock wave interacting with a stationary helium bubble in air that was presented in Section 7.3.2 is considered here. The setup is identical to the one presented in Section 7.3.2, but
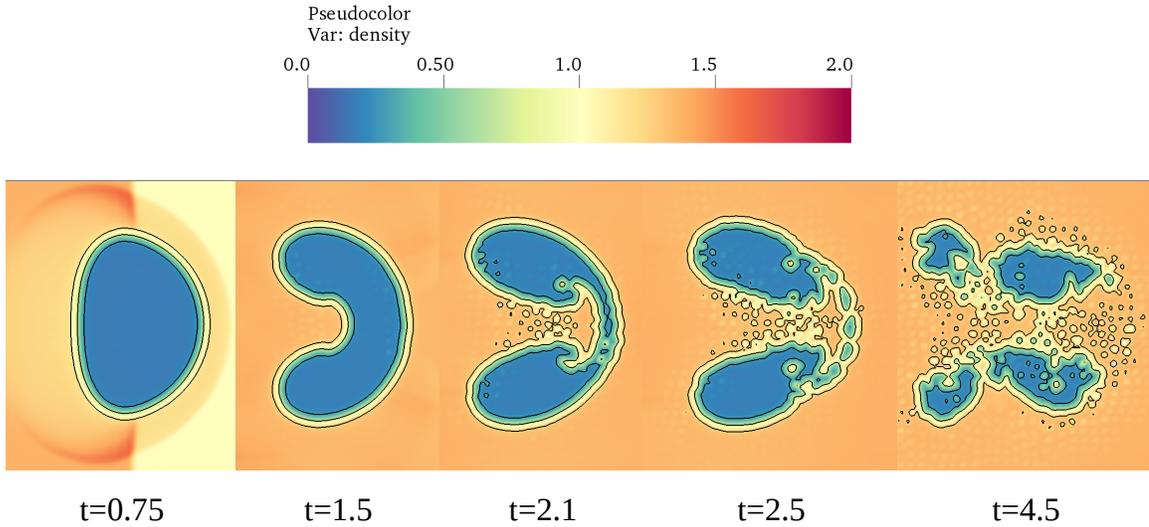
FIGURE 7.12: The bubble shapes at different times for the case of the shock/helium-bubble-in-air interaction using the divergence-form approach, where the interface-regularization terms are discretized using a second-order standard central finite-difference scheme. The three solid black lines denote the isocontours of the volume fraction values of 0.1, 0.5, and 0.9, representing the interface region.

three different grid sizes (600×100, 1200×200, and 2400×400) are considered, where the baseline grid size of 600×100 matches the one used in Section 7.3.2. As the grid is refined, the sharpening coefficient ϵ is reduced so that $\epsilon = \Delta x$. Figure 7.13 shows the snapshots of the interaction of the shock with the bubble on three grids for the LAD, divergence-form, and gradient-form approaches.

For the LAD approach, mixing is still observed in regions of high shear as the grid is refined. However, with grid refinement, the mixing moves to finer scales. On the most refined grid, Kelvin-Helmholtz-type (KH) rollers can be seen all over the bubble interface due to shear between the phases. On the other hand, a constant interface thickness is maintained for the other two approaches. The divergence-form approach leads entrainment of air droplets due to the shear, which become smaller as the grid is refined. The gradient-form approach results in a wavy interface at the locations of unresolved KH rollers.

At the onset of breakup of the bubble at times $t = 2.1$ and $t = 2.5$, a thin film is formed. For the LAD approach, the thin film does not break, but it evolves into a thin region of mixed fluid, which is not physical for two immiscible fluids. However the amount of mixing can be seen to reduce with grid refinement. For the divergence-form approach, the thin film breaks, forming satelitte bubbles due to the lack of grid support. With grid refinement, smaller satelitte bubbles are formed confirming the grid-induced breakup behavior. On the most refined grid, the thin film is seen to be intact, illustrating that the resolution used is sufficient to capture the film at this particular time instant. On the other hand, for the gradient-form approach, the thin film disappears with no formation of
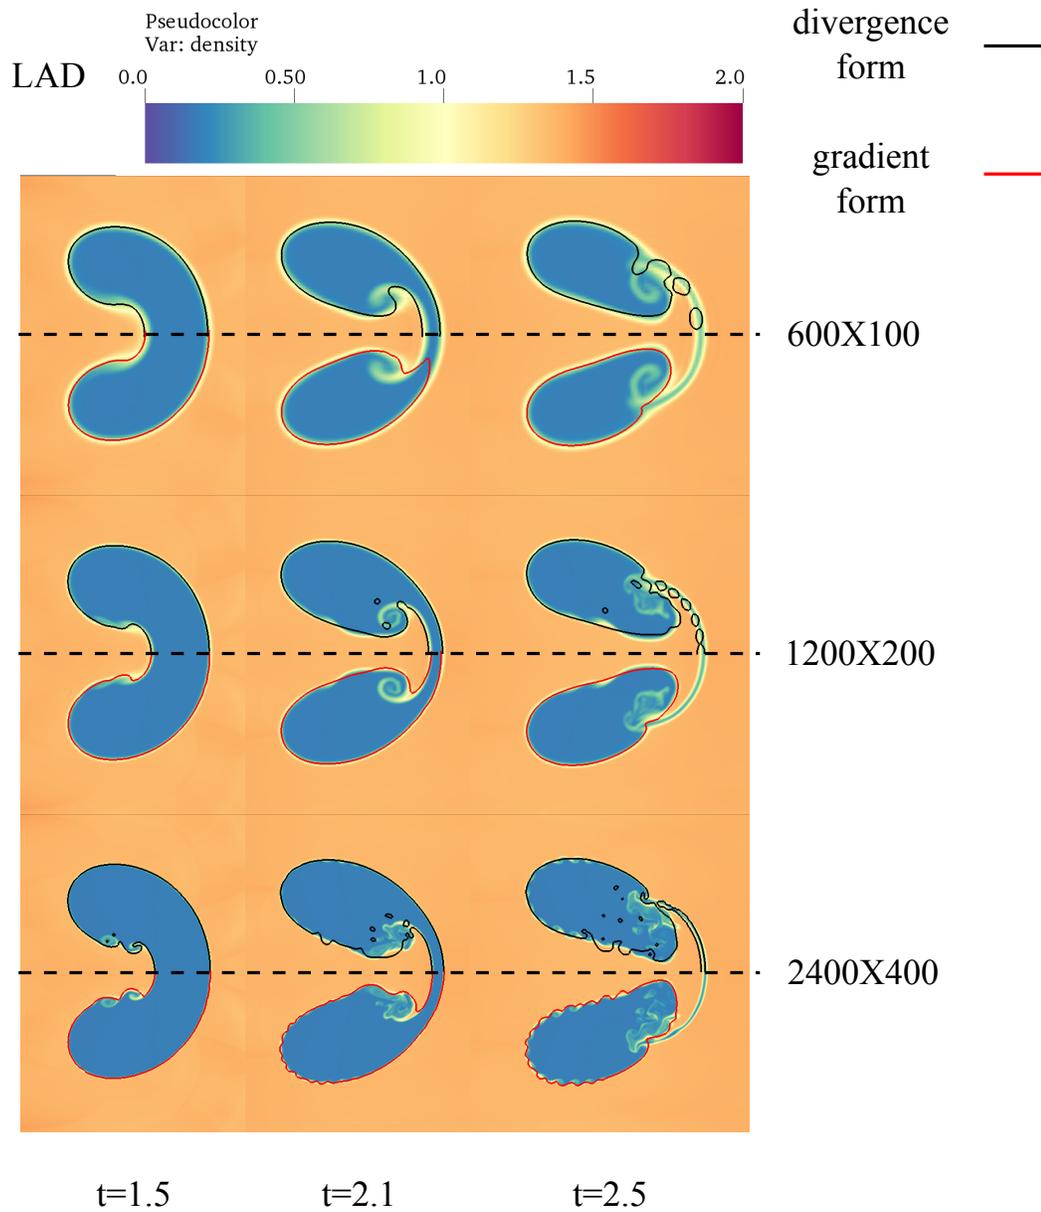
FIGURE 7.13: Comparison of the bubble shapes at different times for the case of the shock/helium-bubble-in-air interaction using various interface-capturing methods and three different grid sizes. The color plot represents the density field obtained using the LAD method, the black solid line on the top half and the red solid line on the bottom half denote the 0.5 isocontour of the volume fraction obtained using the divergence form and the gradient form methods, respectively.
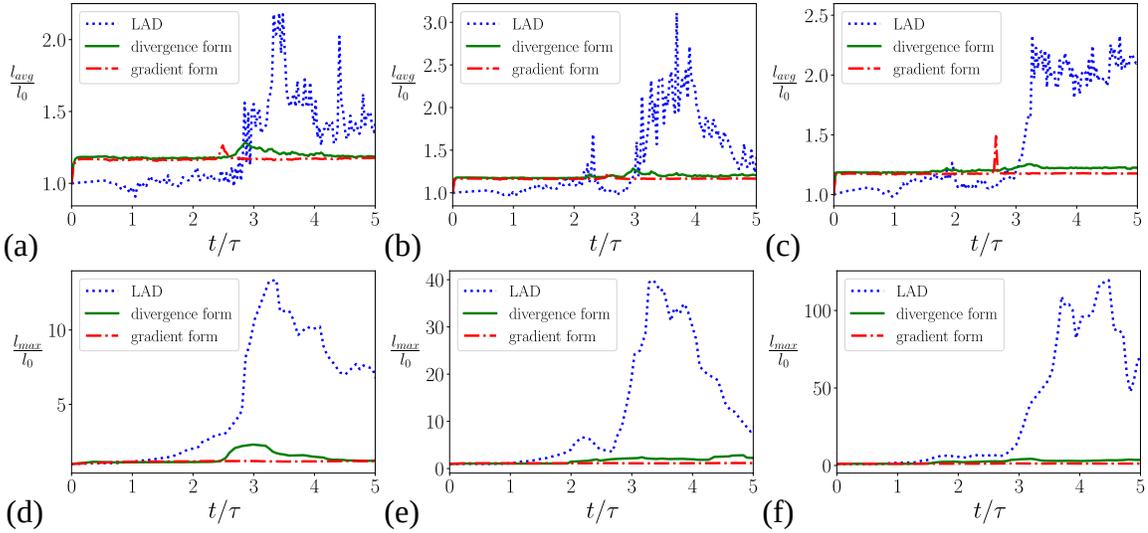
FIGURE 7.14: Comparison of the interface-thickness indicator, l, by various methods and at different grid resolutions, where l_0 is the maximum interface thickness at time $t = 0$. (a), (b), and (c) show the average interface thickness l_{avg} for 600×100, 1200×200, and 2400×400 grids, respectively. (d), (e), and (f) show the maximum interface thickness l_{max} for 600×100, 1200×200, and 2400×400 grids, respectively. To exclude unphysical spikes in l_{max} during breakup events, a moving average of 5 local minima of l_{max} is plotted.

satellite bubbles and the mass is lost due to conservation error. This behavior is persistent with grid refinement at $t = 2.5$, but is reduced with grid refinement at $t = 2.1$, when the bubble is better resolved.

To quantify the changes in interface thickness as the grid is refined, we plot the interface-thickness indicator in Figure 7.14 for the different methods and grid resolutions. Under grid refinement, the normalized average thickness, l_{avg}/l_0, remains qualitatively the same as the baseline resolution of 600×100 for all the methods. This means that, on average, the interface scales with the grid size. On the other hand, the normalized max thickness, l_{max}/l_0, remains the same for the divergence form and the gradient form but increases in magnitude for LAD, due to the widespread fine-scale mixing at high resolution.

To quantify the changes in mass conservation, we plot the mass of the helium bubble in Figure 7.14 for the different methods and grid resolutions. Under grid refinement, the mass conservation properties of the divergence and LAD methods are maintained. At early times, when the interface deformation is still well-resolved, the gradient form method improves in terms of mass conservation, and the three methods agree.

Overall, the result of comparing the three methods on finer grids is that the same differences which were observed at coarse resolution are still observed as the grid is refined. All three methods show improvement in the parts of the flow that are resolved, such as the breaking ligament. In areas
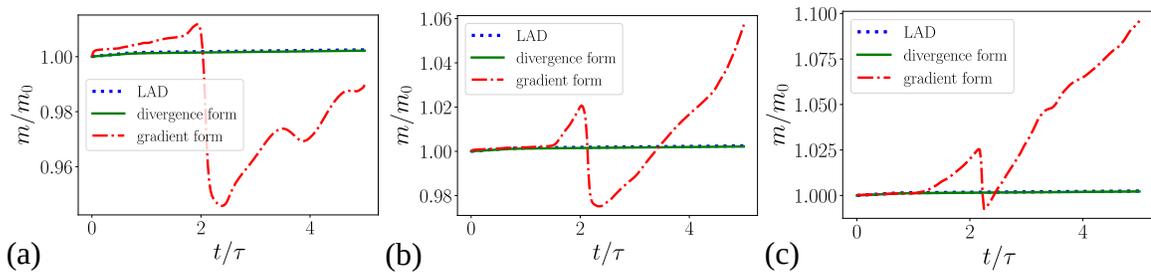
FIGURE 7.15: Comparison of the total mass, m, of the helium bubble by various methods, where m_0 is the mass at time $t = 0$, for (a) the 600×100 grid, (b) the 1200×200 grid, and (c) the 2400×400 grid.

which remained under-resolved, such as the Kelvin-Helmholtz rollers, the methods make the same kinds of errors as before, but they are moved to smaller scales.

Chapter 8

Conclusions

In this thesis, simulations of two different multiphase flows are considered. In the first part (chapters 1-4), high-fidelity simulations of particle-laden channel flow, with and without radiation heating, are performed to increase understanding of the complex physics of this system, which has application to problems of sediment transport, industrial processes, and solar power plants. In the second part (chapters 5-7), a new high-order finite difference method is developed to perform simulations of shock waves in mixtures of elastic-plastic materials undergoing strain hardening due to large, rapid deformation, such as occurs in impact welding and in inertial confinement capsule implosion. A summary of results from each part is presented here, along with directions for future work.

8.1 Flow physics of radiatively heated particle-laden channel flow

To understand the physics of the radiatively heated particle-laden channel flow, a thorough exploration of the fully-developed, isothermal particle-laden channel flow was first conducted. A suite of 10 simulations were performed, covering a wide range of mass loadings (10%-300%) and Stokes numbers ($St^+ = [1, 60]$). These simulations have advanced the state of knowledge about particle-laden flows by using a state-of-the-art velocity correction in the point particle DNS technique to ensure the accuracy of the simulations, as well as examining the variation of Stokes number at a mass loading of 100%, which is higher than previous studies of Stokes number have considered. Using the simulations, the effects of Stokes number and mass loading on the fluid turbulence have been characterized by comparing velocity fluctuations, the balance of stresses, and the budget of turbulent kinetic energy. These simulation results have also been used to test new ideas which represent progress towards turbulence models for two-way coupled particle-laden flows. In particular, by examining the momentum equation in the limit of small Stokes number, a similarity to

variable density turbulence has been identified, and the idea of velocity transforms used in variable density and compressible wall turbulence, has been generalized to particle-laden flows. The mixture transform idea shows significant promise with the "mixture Trettel-Larsson" transform, which collapses particle-laden channel flow velocity data in the viscous sublayer, and collapses the velocity gradients in the logarithmic layer. The transform fails to collapse data in the buffer layer, but the near-constant offset in the log layer is explained by appealing to ideas from the study of drag reduction and rough-wall turbulence: particles disrupt the structure of near-wall eddies, resulting in a thickening of the viscous sublayer. Here, a second modeling idea is introduced by adding an offset into the "mixture Trettel-Larsson" transform. Empirically-tuned offsets can collapse the log layer data, suggesting that the transform has the right mathematical form for this purpose.

The results in Chapter 2 suggest possible avenues for turbulence modeling of two-way coupled particle-laden flows. Previous work on velocity transforms for variable-property flows has been fruitful for the development of improved RANS models [159]. Since the mixture Trettel-Larsson transform developed in the present work uses the same semi-local scaling, similar improvements may be possible for particle-laden wall-bounded flows, provided that an appropriate theory for the roughness-like offsets can be developed. Second, velocity transforms have also proved useful in the development of improved wall models for large-eddy simulation, through an inverse transform [77]. While previous work has shown promise for predicting turbophoresis with wall-modeled LES in one-way coupled flows [103], the more difficult case of two-way coupling has yet to be addressed. The data presented here is a valuable source of information for the effects of near-wall turbulence modulation that must be captured in a wall-modeled LES. In addition, another important task for wall-modeled simulations is to correctly represent the different boundary conditions experienced by the particle and fluid phases, which leads to differences in the velocity fluctuations of fluid and particle phases as the wall is approached.

In terms of testing the validity and usefulness of the transforms developed here, an important step would be to test them on data from other research groups. In particular, it would be important to validate the transforms on data from particle-resolved simulations (such as [45]), which are gradually becoming more affordable, and experimental measurements. Another useful addition to the transform would be accounting for gravity, which has been ignored in this work, but is relevant in most particle-laden flows in the real world.

To complete the characterization of the isothermal particle-laden channel flow, we employed recently-developed approaches based on tessellation to compute the divergence, curl, and helicity of the particle velocity. These quantities in turn represent information about the clustering, rotation, and swirling motions of particle clouds. This study represents the first application of these techniques to channel flow data. The differences in particle motion in the viscous sublayer, buffer layer, and log layers are described. The buffer layer is identified as the region of most intense clustering and vortical motion of particles (greatest variance of divergence and curl) due to the stronger fluctuations in fluid

velocity there. The behavior of particles in the logarithmic layer is found to be similar to previously reported results of one-way coupled particles in homogeneous isotropic turbulence. The quantitative similarity of the dependence of clustering on Stokes number between these flows is striking given the differences in Reynolds number, turbulence anisotropy, and two-way coupling.

Finally, four of the particle-laden flow cases are subjected to strong radiation heating, comparable to what would be expected in a concentrating solar power plant. These simulations cover a regime of higher mass loading and radiation intensity than has been previously studied in simulation or experiment. Radiation heat transfer induces the following changes in the flow: particles absorb radiation and heat up; particles transfer heat convectively to the carrier gas; the carrier gas expands and accelerates due to confinement in the channel; and finally particle clusters expand and accelerate due to momentum coupling with the fluid phase. In addition, because of the turbophoresis effect in the inflow domain, the particle concentration is higher near the walls, and so the particle heat source is localized there. This results in stronger heating (and all of the above effects) near the walls. At larger mass loading, the increasing optical depth also results in radiation attenuation, and so the degree of heating and acceleration decreases going from the bottom wall (where radiation enters the domain) to the top. As these changes take place, there is a sharp transition zone after the radiation boundary condition is applied, during which the near-wall and centerline respond differently. Eventually a quasi-equilibrium is established, and the increase in wall shear stress reaches a new balance with the accelerating flow, which depends on the amount of heat absorbed. As the flow develops, it shows signs of re-laminarizing due to increasing fluid viscosity and the acceleration itself. The scaling ideas developed for the isothermal flow are also tested on the heated flows. The mixture Trettel-Larsson transform shows promise in better collapsing viscous sublayer data across all cases because it accounts for variable fluid density and viscosity. However, because it does not account for flow acceleration, it is not as successful as in the unheated flow. Likewise, the offsets from the unheated flow do not solve the problem of flow acceleration, nor do they account for changes in viscous sublayer thickness as the flow re-laminarizes. The study of the heated flow is concluded by comparing the predictions made with a simple 1D model to bulk quantities computed from the high-fidelity simulation data. Qualitative trends are captured, and the predictions of temperatures, densities, and the predictions for fluid quantities show promising quantitative accuracy.

The analysis of the radiatively heated particle laden channel flow presented here suggests a few directions for future work. First, simulations at higher Reynolds number would help establish whether or not a "quasi-equilibrium" region should be expected in strongly heated particle-laden flows. In the present results, the Reynolds number was low enough that the flow began to relaminarize at the end of the simulation domain, and the quasi-equilibrium seen in the velocity profile began to break down, especially for larger mass loading. Second, conducting simulations at larger optical depth would be important for assessing the accuracy of the simple 1D model in regimes relevant to the solar receiver application. Furthermore, it is reasonable to expect that at large optical depths,

there may be new physical couplings and turbulence-radiation interactions which arise due to strong disparities between the channel walls.

8.2 Simulation methods for shock-driven problems in materials with strength

In this part of thesis, an Eulerian finite difference method for shock-driven deformation in mixtures of elastic-plastic media has been developed. Of particular interest is the regime of large, rapid deformation, in which strain hardening is an important material phenomena, and during which elastic-plastic solids can flow like fluids. An existing code has been extended to treat these new problems by solving evolution equations for the elastic part of the inverse deformation gradient tensor, and for the plastic Finger tensor. New localized artificial diffusion (LAD) terms are introduced into these kinematic equations to improve stability and regularize elastic-plastic shocks. The new diffusivities target shear strain discontinuities, which are a unique characteristic of solids. The LAD framework of Cook [39] was found to be effective at regularizing compressive and tensile discontinuities, whether elastic or plastic, by increasing a tuneable coefficient slightly. In addition to the LAD approach, other modifications were made to the finite difference algorithm to improve robustness in problems with significant shear by discarding rotation information and applying a Gaussian filter to determinant compatibility terms. The new solver has been demonstrated on suite of test cases to demonstrate convergence properties, capability to simulate impact problems, and the interaction of small-scale flow features with shock waves. In addition, a tuning procedure for coefficients in the LAD terms has been developed and described for the various discontinuities which can occur in solids. The solver is also demonstrated on the Taylor impact problem using a Mie-Grüneisen equation of state and a variety of strain hardening models. The simulation results captures the expected deformation behavior, as well as qualitatively capturing the previously reported differences between different empirical plasticity models. Finally, a Richtmyer-Meshkov instability between two solids is simulated to demonstrate the new method's capability with problems involving significant interface distortion and shear deformation. It is shown that strain hardening in this problem arrests the development of the instability, as expected.

In the simulations of the solid-solid Richtmyer-Meshkov instability, especially without strain hardening, a deficiency of the approach described above is that for intense interface roll-up, there is unphysical mixing of solid phases. This occurs because the interface regularization used in the LAD approach is only diffusive. This motivates chapter 7, which considers two additional explicit interface regularization schemes, both of which include a sharpening component: the non-conservative gradient-form approach of Shukla et al. [180] and the conservative divergence-form approach of Jain et al. [97]. These two methods are extended to treat problems involving elastic-plastic solids by deriving compatibility terms which must be added to the evolution equation for the elastic inverse

deformation gradient tensor. All three methods are tested on a variety of problems, including bubble advection, shock-bubble interaction, and the solid-solid Richtmyer-Meshkov instability. All methods perform well when flow features are highly resolved, but differ in their behavior for unresolved flow features. The LAD method causes mixing between the phases; the gradient form induces mass conservation errors, and the divergence form induces interface distortion. In addition, it is demonstrated that if higher-order schemes are used to discretize the sharpening terms, then better shape preservation can be obtained with the divergence form. This was done with a staggered 6th order compact finite differences, and it principle, it would be desirable to extend this to use 10th order compact schemes, to be consistent with how the other terms in the PDE are discretized.

In terms of future work, there are several areas in which the method presented here can be built upon to improve it. One way is to improve upon the formally first-order temporal convergence. To do this first requires a higher order scheme for the implicit plastic relaxation, This could in principle be obtained by applying an implicit-explicit Runge-Kutta scheme [156], or by using a recently developed technique to pre-condition the stiff plastic relaxation term with desired order of accuracy for time integration [22]. The second requirement to achieve higher-order temporal convergence is to improve upon the rotation-treatment procedure, which also limits to first-order at present.

Another area of potential improvement is in the localization of the diffusivities in the kinematic equations. In the present method, the LAD is localized to ringing in a norm of the elastic or plastic strain tensors, and a switching function is further used to concentrate the diffusion on areas dominated by vorticity, but this approach is not perfect. For example, if a material undergoing extreme shear has a shock wave pass through it, then it will receive less attention from the LAD, which is undesirable. This could be improved by moving away from the switching function approach and instead changing to a different ringing detector. Instead of the strain norm used a now, a different quantity, based on tensor invariants which are sensitive to shear but agnostic to volume change, could be formulated to appropriately add diffusion where there is shear, regardless of the dilatation.

Bibliography

[1] M. Abdelrahman, P. Fumeaux, and P. Suter. Study of solid-gas-suspensions used for direct absorption of concentrated solar radiation. *Solar Energy*, 22(1):45–48, 1979. ISSN 0038-092X. doi: https://doi.org/10.1016/0038-092X(79)90058-6.

[2] M.C. Adler, J.R. West, and S.K. Lele. Differential operators for high-rank tensors in cylindrical and spherical coordinates. *Center for Turbulence Research Annual Research Briefs*, pages 231–242, 2021.

[3] Michael C Adler and Datta V Gaitonde. Dynamic linear response of a shock/turbulent-boundary-layer interaction using constrained perturbations. *Journal of Fluid Mechanics*, 840:291–341, 2018. doi: doi:10.1017/jfm.2018.70.

[4] Michael C Adler and Datta V Gaitonde. Dynamics of strong swept-shock/turbulent-boundary-layer interactions. *Journal of Fluid Mechanics*, 2020. In Press.

[5] Michael C. Adler and Sanjiva K. Lele. Strain-hardening framework for Eulerian simulations of multi-material elasto-plastic deformation. *Center for Turbulence Research Annual Research Briefs*, pages 257–271, 2019.

[6] Grégoire Allaire, Sébastien Clerc, and Samuel Kokh. A five-equation model for the simulation of interfaces between compressible fluids. *Journal of Computational Physics*, 181(2):577–616, 2002.

[7] H. Andersson, L. Zhao, and E. Variano. On the anisotropic vorticity in turbulent channel flows. *J. Fluids Eng*, 137:084503, 2015.

[8] MR Ansari and A Daramizadeh. Numerical simulation of compressible two-phase flow using a diffuse interface method. *International Journal of Heat and Fluid Flow*, 42:209–223, 2013.

[9] Mohamad Aslani and Jonathan D Regele. A localized artificial diffusivity method to simulate compressible multiphase flows using the stiffened gas equation of state. *International Journal for Numerical Methods in Fluids*, 88(9):413–433, 2018.

[10] F. Aurenhammer. Voronoi diagrams—a survey of a fundamental geometric data structure. *ACM Comput. Surv.*, 23:345–405, 1991.

[11] MR Baer and JW Nunziato. A two-phase mixture theory for the deflagration-to-detonation transition (DDT) in reactive granular materials. *International journal of multiphase flow*, 12 (6):861–889, 1986.

[12] Xiao Bai and Xiaolong Deng. A sharp interface method for compressible multi-phase flows based on the cut cell and ghost fluid methods. *Advances in Applied Mathematics and Mechanics*, 9(5):1052–1075, 2017.

[13] S. Balachandar and J. Eaton. Turbulent dispersed multiphase flow. *Annual Review Of Fluid Mechanics*, 42(1):111–133, 2010. doi: 10.1146/annurev.fluid.010908.165243.

[14] Biswajit Banerjee. An evaluation of plastic flow stress models for the simulation of high-temperature and high-strain-rate deformation of metals. *arXiv preprint cond-mat/0512466*, 2005.

[15] S. Banerjee, R. Krahl, F. Durst, and Ch. Zenger. Presentation of anisotropy properties of turbulence, invariants versus eigenvalue approaches. *Journal of Turbulence*, 8:N32, 2007. doi: 10.1080/14685240701506896.

[16] Andrew J. Banko, Laura Villafañe, Ji Hoon Kim, and John K. Eaton. Temperature statistics in a radiatively heated particle-laden turbulent square duct flow. *International Journal of Heat and Fluid Flow*, 84:108618, 2020. ISSN 0142-727X. doi: https://doi.org/10.1016/j.ijheatfluidflow.2020.108618.

[17] Andrew James Banko. *Radiation absorption by inertial particles in a turbulent square duct flow*. PhD thesis, Stanford University, 2018.

[18] C. Barber, D. Dobkin, and H. Huhdanpaa. The quickhull algorithm for convex hulls. *ACM Trans. Math. Softw.*, 22:469–483, 1996.

[19] Philip Barton and Evgeniy Romenski. On computational modelling of strain-hardening material dynamics. *Communications in Computational Physics*, 11(5):1525–1546, 2012.

[20] Philip T. Barton. An interface-capturing godunov method for the simulation of compressible solid-fluid problems. *Journal of Computational Physics*, 390:25–50, 2019. ISSN 0021-9991. doi: https://doi.org/10.1016/j.jcp.2019.03.044.

[21] Philip T Barton, Ralf Deiterding, Daniel Meiron, and Dale Pullin. Eulerian adaptive finite-difference method for high-velocity impact and penetration problems. *Journal of Computational Physics*, 240:76–99, 2013.

[22] Maxime Bassenne, Lin Fu, and Ali Mani. Time-accurate and highly-stable explicit operators for stiff differential equations. *Journal of Computational Physics*, 424:109847, 2021. ISSN 0021-9991. doi: https://doi.org/10.1016/j.jcp.2020.109847.

[23] P. Basu. *Combustion and Gasification in Fluidized Beds*. CRC Press, 2006.

[24] F. Battista, J.-P. Mollicone, P. Gualtieri, R. Messina, and C. M. Casciola. Exact regularised point particle (ERPP) method for particle-laden wall-bounded flows in the two-way coupling regime. *Journal of Fluid Mechanics*, 878:420–444, 2019. doi: 10.1017/jfm.2019.622.

[25] Shahaboddin Alahyari Beig and Eric Johnsen. Maintaining interface equilibrium conditions in compressible multiphase flows using interface capturing. *Journal of Computational Physics*, 302:548–566, 2015.

[26] David J Benson. Computational methods in Lagrangian and Eulerian hydrocodes. *Computer Methods in Applied Mechanics and Engineering*, 99(2-3):235–394, 1992.

[27] Wurigen Bo and John W Grove. A volume of fluid method based ghost fluid method for compressible multi-fluid flows. *Computers & Fluids*, 90:113–122, 2014.

[28] S. Bose. *Explicitly Filtered Large-Eddy Simulation: With Application to Grid Adaptation and Wall Modeling*. PhD thesis, Stanford University, 2012.

[29] NK Bourne and JE Field. Shock-induced collapse of single cavities in liquids. *Journal of Fluid Mechanics*, 244:225–240, 1992.

[30] A. Bragg, D. Richter, and G. Wang. Settling strongly modifies particle concentrations in wall-bounded turbulent flows even when the settling parameter is asymptotically small. *Phys. Rev. Fluids*, 6:124301, 2021.

[31] L. Brandt and F. Coletti. Particle-laden turbulence: Progress and perspectives. *Annual Review Of Fluid Mechanics*, 54:159–189, 2022.

[32] Christopher E Brennen. *Fundamentals of multiphase flow*. Cambridge university press, 2005.

[33] Jesse Capecelatro and Olivier Desjardins. Mass Loading Effects on Turbulence Modulation by Particle Clustering in Dilute and Moderately Dilute Channel Flows. *Journal of Fluids Engineering*, 137(11), 11 2015. ISSN 0098-2202. doi: 10.1115/1.4030644. 111102.

[34] Jesse Capecelatro, Olivier Desjardins, and Rodney O. Fox. Strongly coupled fluid-particle flows in vertical channels. I. Reynolds-averaged two-phase turbulence statistics. *Physics of Fluids*, 28(3):033306, 2016. doi: 10.1063/1.4943231.

[35] Jesse Capecelatro, Olivier Desjardins, and Rodney O. Fox. On the transition between turbulence regimes in particle-laden channel flows. *Journal of Fluid Mechanics*, 845:499–519, 2018. doi: 10.1017/jfm.2018.259.

[36] Marion Capuano, Christophe Bogey, and Peter DM Spelt. Simulations of viscous and compressible gas–gas flows using high-order finite difference schemes. *Journal of Computational Physics*, 361:56–81, 2018.

[37] Alexandre Chiapolino, Richard Saurel, and Boniface Nkonga. Sharpening diffuse interfaces with compressible fluids on unstructured meshes. *Journal of Computational Physics*, 340: 389–417, 2017.

[38] Robert Chiodi and Olivier Desjardins. A reformulation of the conservative level set reinitialization equation for accurate and robust simulation of complex multiphase flows. *Journal of Computational Physics*, 343:186–200, 2017.

[39] Andrew W Cook. Artificial fluid properties for large-eddy simulation of compressible turbulent mixing. *Physics of Fluids*, 19(5):055103, 2007.

[40] Andrew W Cook. Enthalpy diffusion in multicomponent flows. *Physics of Fluids*, 21(5):055109, 2009.

[41] Andrew W Cook and William H Cabot. Hyperviscosity for shock-turbulence interactions. *Journal of Computational Physics*, 203(2):379–385, 2005.

[42] Vedran Coralic and Tim Colonius. Finite-volume WENO scheme for viscous compressible multicomponent flows. *Journal of Computational Physics*, 274:95–121, 2014.

[43] Pedro Costa, Francesco Picano, Luca Brandt, and Wim-Paul Breugem. Universal scaling laws for dense particle suspensions in turbulent wall-bounded flows. *Phys. Rev. Lett.*, 117:134501, Sep 2016. doi: 10.1103/PhysRevLett.117.134501.

[44] Pedro Costa, Luca Brandt, and Francesco Picano. Interface-resolved simulations of small inertial particles in turbulent channel flow. *Journal of Fluid Mechanics*, 883:A54, 2020. doi: 10.1017/jfm.2019.918.

[45] Pedro Costa, Luca Brandt, and Francesco Picano. Near-wall turbulence modulation by small inertial particles. *Journal of Fluid Mechanics*, 922:A9, 2021. doi: 10.1017/jfm.2021.507.

[46] R. S. Craxton, K. S. Anderson, T. R. Boehly, V. N. Goncharov, D. R. Harding, J. P. Knauer, R. L. McCrory, P. W. McKenty, D. D. Meyerhofer, J. F. Myatt, A. J. Schmitt, J. D. Sethian, R. W. Short, S. Skupsky, W. Theobald, W. L. Kruer, K. Tanaka, R. Betti, T. J. B. Collins, J. A. Delettrez, S. X. Hu, J. A. Marozas, A. V. Maximov, D. T. Michel, P. B. Radha, S. P.

Regan, T. C. Sangster, W. Seka, A. A. Solodov, J. M. Soures, C. Stoeckl, and J. D. Zuegel. Direct-drive inertial confinement fusion: A review. *Physics of Plasmas*, 22(11):110501, 2015. doi: 10.1063/1.4934714.

[47] Frederic Daude, Pascal Galon, Zhenlan Gao, and Etienne Blaud. Numerical experiments using a HLLC-type scheme with ALE formulation for compressible two-phase flows five-equation models with phase transition. *Computers & Fluids*, 94:112–138, 2014.

[48] Alban de Vaucorbeil, Vinh Phu Nguyen, and Christopher R. Hutchinson. A total-lagrangian material point method for solid mechanics problems involving large deformations. *Computer Methods in Applied Mechanics and Engineering*, 360:112783, 2020. ISSN 0045-7825. doi: https://doi.org/10.1016/j.cma.2019.112783.

[49] N.G. Deen, M. Van Sint Annaland, M.A. Van der Hoef, and J.A.M. Kuipers. Review of discrete particle modeling of fluidized beds. *Chemical Engineering Science*, 62(1):28–44, 2007. ISSN 0009-2509. doi: https://doi.org/10.1016/j.ces.2006.08.014. Fluidized Bed Applications.

[50] Xi Deng, Satoshi Inaba, Bin Xie, Keh-Ming Shyue, and Feng Xiao. High fidelity discontinuity-resolving reconstruction for compressible multiphase flows with moving interfaces. *Journal of Computational Physics*, 371:945–966, 2018.

[51] Guy Dimonte, G. Terrones, F. J. Cherne, T. C. Germann, V. Dupont, K. Kadau, W. T. Buttler, D. M. Oro, C. Morris, and D. L. Preston. Use of the Richtmyer-Meshkov instability to infer yield stress at high-energy densities. *Phys. Rev. Lett.*, 107:264502, Dec 2011. doi: 10.1103/PhysRevLett.107.264502.

[52] Jean Donea, Antonio Huerta, Jean-Philippe Ponthot, and Antonio Rodríguez-Ferran. Arbitrary Lagrangian–Eulerian methods. *Encyclopedia of Computational Mechanics*, 1:413–437, 2004. Chapter 14.

[53] M. Esmaily and J. Horwitz. A correction scheme for two-way coupled point-particle simulations on anisotropic grids. *J. Comput. Phys.*, 375:960–982, 2018. ISSN 0021-9991. doi: https://doi.org/10.1016/j.jcp.2018.09.009.

[54] M. Esmaily, L. Jofre, A. Mani, and G. Iaccarino. A scalable geometric multigrid solver for non-symmetric elliptic systems with application to variable-density flows. *Journal of Computational Physics*, 357:142–158, 2018. ISSN 0021-9991. doi: https://doi.org/10.1016/j.jcp.2017.12.024.

[55] M. Esmaily, L. Villafane, A.J. Banko, G. Iaccarino, J.K. Eaton, and A. Mani. A benchmark for particle-laden turbulent duct flow: A joint computational and experimental study. *International Journal of Multiphase Flow*, 132:103410, 2020. ISSN 0301-9322. doi: https://doi.org/10.1016/j.ijmultiphaseflow.2020.103410.

[56] Mahdi Esmaily and Ali Mani. Modal analysis of the behavior of inertial particles in turbulence subjected to stokes drag. *Phys. Rev. Fluids*, 5:084303, Aug 2020. doi: 10.1103/PhysRevFluids. 5.084303.

[57] Mahdi Esmaily-Moghadam and Ali Mani. Analysis of the clustering of inertial particles in turbulent flows. *Phys. Rev. Fluids*, 1:084202, Dec 2016. doi: 10.1103/PhysRevFluids.1.084202.

[58] M. Farge, G. Pellegrino, and K. Schneider. Coherent vortex extraction in 3d turbulent flows using orthogonal wavelets. *Physical Review Letters*, 87(5):054501, 2001.

[59] Nicolas Favrie and Sergey Gavrilyuk. Mathematical and numerical model for nonlinear viscoplasticity. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 369(1947):2864–2880, 2011.

[60] Nicolas Favrie and Sergey L Gavrilyuk. Diffuse interface model for compressible fluid–compressible elastic–plastic solid interaction. *Journal of Computational Physics*, 231(7):2695–2723, 2012.

[61] Ronald P Fedkiw, Tariq Aslam, Barry Merriman, Stanley Osher, et al. A non-oscillatory Eulerian approach to interfaces in multimaterial flows (the ghost fluid method). *Journal of Computational Physics*, 152(2):457–492, 1999.

[62] W. A. Fiveland. Discrete-Ordinates Solutions of the Radiative Transport Equation for Rectangular Enclosures. *Journal of Heat Transfer*, 106(4):699–706, 11 1984. ISSN 0022-1481. doi: 10.1115/1.3246741.

[63] Marianne M Francois, Sharen J Cummins, Edward D Dendy, Douglas B Kothe, James M Sicilian, and Matthew W Williams. A balanced-force algorithm for continuous and sharp interfacial surface tension models within a volume tracking framework. *Journal of Computational Physics*, 213(1):141–173, 2006.

[64] A. Frankel, G. Iaccarino, and A. Mani. Optical depth in particle-laden turbulent flows. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 201:10–16, 2017. ISSN 0022-4073. doi: https://doi.org/10.1016/j.jqsrt.2017.06.029.

[65] Ari Frankel. *Modeling radiation transport in turbulent particle-laden media*. PhD thesis, Stanford University, 2017.

[66] Ari Frankel, H. Pouransari, F. Coletti, and A. Mani. Settling of heated particles in homogeneous turbulence. *Journal of Fluid Mechanics*, 792:869–893, 2016. doi: 10.1017/jfm.2016.102.

[67] Pierre Février, Olivier Simonin, and Kyle D. Squires. Partitioning of particle velocities in gas–solid turbulent flows into a continuous field and a spatially uncorrelated random distribution: theoretical formalism and numerical study. *Journal of Fluid Mechanics*, 533:1–46, 2005. doi: 10.1017/S0022112005004088.

[68] Swetava Ganguli and Sanjiva K. Lele. Drag of a heated sphere at low reynolds numbers in the absence of buoyancy. *Journal of Fluid Mechanics*, 869:264–291, 2019. doi: 10.1017/jfm.2019. 187.

[69] Daniel P Garrick, Wyatt A Hagen, and Jonathan D Regele. An interface capturing scheme for modeling atomization in compressible flows. *Journal of Computational Physics*, 344:260–280, 2017.

[70] Daniel P Garrick, Mark Owkes, and Jonathan D Regele. A finite-volume HLLC-based scheme for compressible interfacial flows with surface tension. *Journal of Computational Physics*, 339: 46–67, 2017.

[71] Sergey L Gavrilyuk, Nicolas Favrie, and Richard Saurel. Modelling wave dynamics of compressible elastic materials. *Journal of Computational Physics*, 227(5):2941–2969, 2008.

[72] N S Ghaisas, A Subramaniam, S K Lele, and A W Cook. Evaluation of an Eulerian multimaterial mixture formulation based on a single inverse deformation gradient tensor field. *Center for Turbulence Research Annual Briefs*, 2017(LLNL-JRNL-741479), 2017.

[73] Niranjan S Ghaisas, Akshay Subramaniam, and Sanjiva K Lele. A unified high-order Eulerian method for continuum simulations of fluid flow and of elastic–plastic deformations in solids. *Journal of Computational Physics*, 371:452–482, 2018.

[74] A S Ghate, A Subramaniam, N Ghaisas, J R West, M C Adler, and S S Jain. PadéOps source code, 2021. DOI: 10.5281/zenodo.5277989.

[75] S K Godunov and E I Romenskii. Nonstationary equations of nonlinear elasticity theory in Eulerian coordinates. *Journal of Applied Mechanics and Technical Physics*, 13(6):868–884, 1972.

[76] Giorgio Greto and Sivakumar Kulasegaram. An efficient and stabilised sph method for large strain metal plastic deformations. *Computational Particle Mechanics*, 7(3):523–539, May 2020. ISSN 2196-4386. doi: 10.1007/s40571-019-00277-6.

[77] Kevin Griffin. *Modeling and simulation of wall-bounded turbulent flows with pressure-gradient and compressibility effects.* PhD thesis, Stanford University, 2022.

[78] Kevin Patrick Griffin, Lin Fu, and Parviz Moin. Velocity transformation for compressible wall-bounded turbulent flows with and without heat transfer. *Proceedings of the National Academy of Sciences*, 118(34), 2021. ISSN 0027-8424. doi: 10.1073/pnas.2111144118.

[79] J F Haas and Bradford Sturtevant. Interaction of weak shock waves with cylindrical and spherical gas inhomogeneities. *Journal of Fluid Mechanics*, 181:41–76, 1987.

[80] Ory Haimovich and Steven H Frankel. Numerical simulations of compressible multicomponent and multiphase flow using a high-order targeted ENO (TENO) finite-volume method. *Computers & Fluids*, 146:105–116, 2017.

[81] Sarah Hank, Sergey Gavrilyuk, Nicolas Favrie, and Jacques Massoni. Impact simulation by an Eulerian model for interaction of multiple elastic–plastic solids and fluids. *International Journal of Impact Engineering*, 109:104–111, 2017.

[82] David J Hill, D Pullin, Michael Ortiz, and D Meiron. An Eulerian hybrid WENO centered-difference solver for elastic–plastic solids. *Journal of Computational Physics*, 229(24):9053–9072, 2010.

[83] C.W. Hirt, A.A. Amsden, and J.L. Cook. An arbitrary lagrangian–eulerian computing method for all flow speeds. *Journal of Computational Physics*, 135(2):203–216, 1997. ISSN 0021-9991. doi: https://doi.org/10.1006/jcph.1997.5702.

[84] Clifford K. Ho. A review of high-temperature particle receivers for concentrating solar power. *Applied Thermal Engineering*, 109:958–969, 2016. ISSN 1359-4311. doi: https://doi.org/10.1016/j.applthermaleng.2016.04.103. Special Issue: Solar Energy Research Institute for India and the United States (SERIIUS) – Concentrated Solar Power.

[85] Clifford K. Ho. Advances in central receivers for concentrating solar applications. *Solar Energy*, 152:38–56, 2017. ISSN 0038-092X. doi: https://doi.org/10.1016/j.solener.2017.03.048. Progress in Solar Energy Special Issue: Concentrating Solar Power (CSP).

[86] Clifford K. Ho and Brian D. Iverson. Review of high-temperature central receiver designs for concentrating solar power. *Renewable and Sustainable Energy Reviews*, 29:835–846, 2014. ISSN 1364-0321. doi: https://doi.org/10.1016/j.rser.2013.08.099.

[87] J.A.K. Horwitz and A. Mani. Accurate calculation of Stokes drag for point–particle tracking in two-way coupled flows. *Journal of Computational Physics*, 318:85–109, 2016. ISSN 0021-9991. doi: https://doi.org/10.1016/j.jcp.2016.04.034.

[88] Jeremy A. K. Horwitz, Swetava Ganguli, Sanjiva K. Lele, and Ali Mani. Settling of two-way momentum and energy coupled particles subject to Boussinesq and non-Boussinesq heating.

Theoretical and Computational Fluid Dynamics, 35(4):539–551, Aug 2021. ISSN 1432-2250. doi: 10.1007/s00162-021-00572-0.

[89] Xiang Yu Hu and Boo Cheong Khoo. An interface interaction method for compressible multifluids. *Journal of Computational Physics*, 198(1):35–64, 2004.

[90] A J Hunt. New solar thermal receiver utilizing a small particle heat exchanger. Technical report, Lawrence Berkeley National Laboratory, 4 1979.

[91] K. Hutter. Geophysical granular and particle-laden flows: review of the field. *Philos. Trans. R. Soc. A*, 363:1497–1505, 2005.

[92] Suhas S Jain. Accurate conservative phase-field method for simulation of two-phase flows. *Journal of Computational Physics*, 469:111529, 2022.

[93] Suhas S Jain and Parviz Moin. A kinetic energy and entropy preserving scheme for the simulation of compressible two-phase turbulent flows. *Center for Turbulence Research Annual Research Briefs*, 2020.

[94] Suhas S Jain and Parviz Moin. A kinetic energy–and entropy-preserving scheme for compressible two-phase flows. *Journal of Computational Physics*, page 111307, 2022.

[95] Suhas S Jain, Ali Mani, and Parviz Moin. A conservative diffuse-interface method for the simulation of compressible two-phase flows with turbulence and acoustics. *Center for Turbulence Research Annual Research Briefs*, pages 47–64, 2018.

[96] Suhas S Jain, Ken Kamrin, and Ali Mani. A conservative and non-dissipative Eulerian formulation for the simulation of soft solids in fluids. *Journal of Computational Physics*, 399: 108922, 2019.

[97] Suhas S Jain, Ali Mani, and Parviz Moin. A conservative diffuse-interface method for compressible two-phase flows. *Journal of Computational Physics*, page 109606, 2020.

[98] Suhas S. Jain, Michael C. Adler, Jacob R. West, Ali Mani, Parviz Moin, and Sanjiva K. Lele. Assessment of diffuse-interface methods for compressible multiphase fluid flows and elastic-plastic deformation in solids. *Journal of Computational Physics*, page 111866, 2022. ISSN 0021-9991. doi: https://doi.org/10.1016/j.jcp.2022.111866.

[99] Eric Johnsen and Tim Colonius. Implementation of WENO schemes in compressible multi-component flow problems. *Journal of Computational Physics*, 219(2):715–732, 2006.

[100] Eric Johnsen and Frank Ham. Preventing numerical errors generated by interface-capturing schemes in compressible multi-material flows. *Journal of Computational Physics*, 231(17): 5705–5717, 2012.

[101] Eric Johnsen, Johan Larsson, Ankit V Bhagatwala, William H Cabot, Parviz Moin, Britton J Olson, Pradeep S Rawat, Santhosh K Shankar, Björn Sjögreen, H C Yee, et al. Assessment of high-resolution methods for numerical simulations of compressible turbulence with shock waves. *Journal of Computational Physics*, 229(4):1213–1237, 2010.

[102] G.R. Johnson and W.H. Cook. A constitutive model and data for metals subjected to large strain rates and high temperatures. In *Proceedings of the Seventh International Symposium on Ballistics*, pages 541–547, The Hague, The Netherlands, 1983.

[103] Perry L. Johnson, Maxime Bassenne, and Parviz Moin. Turbophoresis of small inertial particles: theoretical considerations and application to wall-modelled large-eddy simulations. *Journal of Fluid Mechanics*, 883:A27, 2020. doi: 10.1017/jfm.2019.865.

[104] Ken Kamrin, Chris H. Rycroft, and Jean-Christophe Nave. Reference map technique for finite-strain elasticity and fluid–solid interaction. *Journal of the Mechanics and Physics of Solids*, 60(11):1952–1969, 2012. ISSN 0022-5096. doi: https://doi.org/10.1016/j.jmps.2012.06.003.

[105] A K Kapila, R Menikoff, J B Bdzil, S F Son, and D Scott Stewart. Two-phase modeling of deflagration-to-detonation transition in granular materials: Reduced equations. *Physics of Fluids*, 13(10):3002–3024, 2001.

[106] Isao Kataoka. Local instant formulation of two-phase flow. *International Journal of Multiphase Flow*, 12(5):745–758, 1986.

[107] S. Kawai, S. K. Shankar, and S. K. Lele. Assessment of localized artificial diffusivity scheme for large-eddy simulation of compressible turbulent flows. *Journal of Computational Physics*, 229(5):1739–1762, 2010.

[108] Soshi Kawai and Sanjiva K Lele. Localized artificial diffusivity scheme for discontinuity capturing on curvilinear meshes. *Journal of Computational Physics*, 227(22):9498–9526, 2008.

[109] Christopher A Kennedy, Mark H Carpenter, and R Michael Lewis. Low-storage, explicit Runge–Kutta schemes for the compressible Navier–Stokes equations. *Applied numerical mathematics*, 35(3):177–219, 2000.

[110] John Kim, Parviz Moin, and Robert Moser. Turbulence statistics in fully developed channel flow at low reynolds number. *Journal of Fluid Mechanics*, 177:133–166, 1987. doi: 10.1017/S0022112087000892.

[111] Donald L Koch and Reghan J Hill. Inertial effects in suspension and porous-media flows. *Annual Review of Fluid Mechanics*, 33(1):619–647, 2001. doi: 10.1146/annurev.fluid.33.1.619.

[112] J. G. M. Kuerten. Point-particle DNS and and LES of particle-laden turbulent flow-a state-of-the-art review. *Flow, Turbulence and Combustion*, 97:689–713, 2016.

[113] J. D. Kulick, J. R. Fessler, and J. K. Eaton. Particle response and turbulence modification in fully developed channel flow. *Journal of Fluid Mechanics*, 277:109–134, 1994. doi: 10.1017/S0022112094002703.

[114] J. Lee and C. Lee. Modification of particle-laden near-wall turbulence: Effect of Stokes number. *Physics of Fluids*, 27(2):023303, 2015. doi: 10.1063/1.4908277.

[115] J. Lee and C. Lee. The effect of wall-normal gravity on particle-laden near-wall turbulence. *Journal of Fluid Mechanics*, 873:475–507, 2019.

[116] Sanjiva K Lele. Compact finite difference schemes with spectral-like resolution. *Journal of Computational Physics*, 103(1):16–42, 1992.

[117] Yiming Li, J. B. McLaughlin, K. Kontomaris, and L. Portela. Numerical simulation of particle-laden turbulent channel flow. *Physics of Fluids*, 13(10):2957–2967, 2001. doi: 10.1063/1.1396846.

[118] John Lindl, Otto Landen, John Edwards, Ed Moses, and NIC Team. Review of the National Ignition Campaign 2009-2012. *Physics of Plasmas*, 21(2):020501, 2014.

[119] Kai Liu, Mandar Lakhote, and S. Balachandar. Self-induced temperature correction for interphase heat transfer in Euler-Lagrange point-particle simulation. *Journal of Computational Physics*, 396:596–615, 2019. ISSN 0021-9991. doi: https://doi.org/10.1016/j.jcp.2019.06.069.

[120] Alejandro Lopez Ortega. *Simulation of Richtmyer-Meshkov flows for elastic-plastic solids in planar and converging geometries using an Eulerian framework*. PhD thesis, California Institute of Technology, 2013.

[121] Andrew Majda and James Sethian. The derivation and numerical solution of the equations for zero Mach number combustion. *Combustion Science and Technology*, 42(3-4):185–205, 1985. doi: 10.1080/00102208508960376.

[122] Ali Mani, Johan Larsson, and Parviz Moin. Suitability of artificial bulk viscosity for large-eddy simulation of turbulent flows with shocks. *Journal of Computational Physics*, 228(19):7368–7374, 2009.

[123] C. Marchioli and A. Soldati. Mechanisms for particle transfer and segregation in a turbulent boundary layer. *J. Fluid Mech.*, 468:283–315, 2002.

[124] Cristian Marchioli. Large-eddy simulation of turbulent dispersed flows: a review of modelling approaches. *Acta Mechanica*, 228(3):741–771, 2017.

[125] Antonio Marquina and Pep Mulet. A flux-split algorithm applied to conservative models for multicomponent compressible flows. *Journal of Computational Physics*, 185(1):120–138, 2003.

[126] K. Matsuda, R. Onishi, R. Kurose, and S. Komori. Turbulence effect on cloud radiation. *Phys. Rev. Lett.*, 108:224502, 2012.

[127] K. Matsuda, R. Onishi, M. Hirahara, R. Kurose, K. Takahashi, and S. Komori. Influence of microscale turbulent droplet clustering on radar cloud observations. *J. Atmos. Sci*, 71: 3569–3582, 2014.

[128] K. Matsuda, K. Schneider, T. Oujia, J. West, S. Jain, and K. Maeda. Multiresolution analysis of inertial particle tessellations for clustering dynamics. In *Proceedings of the Summer Program*, Stanford University, 2022. Center for Turbulence Research.

[129] M. R. Maxey. The gravitational settling of aerosol particles in homogeneous turbulence and random flow fields. *Journal of Fluid Mechanics*, 174:441–465, 1987. doi: 10.1017/ S0022112087000193.

[130] Martin R. Maxey and James J. Riley. Equation of motion for a small rigid sphere in a nonuniform flow. *The Physics of Fluids*, 26(4):883–889, 1983. doi: 10.1063/1.864230.

[131] Jure Mencinger and Iztok Žun. On the finite volume discretization of discontinuous body force field on collocated grid: Application to VOF method. *Journal of Computational Physics*, 221 (2):524–538, 2007.

[132] R Menikoff. Elastic-plastic shock waves. In Y Horie, editor, *Shock Wave Science and Technology Reference Library Volume 2*, pages 189–224. Springer, Berlin, Germany, 2007.

[133] LA Merzhievskii and AD Resnyanskii. Numerical analysis for the abnormality of protective properties of a thin shield under oblique impact. *Combustion, Explosion and Shock Waves*, 29 (6):744–749, 1993.

[134] Fletcher J Miller and Roland W Koenigsdorff. Thermal modeling of a small-particle solar central receiver. *J. Sol. Energy Eng.*, 122(1):23–29, 2000.

[135] G H Miller and P Colella. A high-order Eulerian Godunov method for elastic–plastic flow in solids. *Journal of Computational Physics*, 167(1):131–176, 2001.

[136] Shahab Mirjalili, Suhas S Jain, and Micheal Dodd. Interface-capturing methods for two-phase flows: An overview and recent developments. *Center for Turbulence Research Annual Research Briefs*, pages 117–135, 2017.

[137] H. K. Moffatt and A. Tsinober. Helicity in laminar and turbulent flow. *Annual Review of Fluid Mechanics*, 24(1):281–312, 1992.

[138] R. Monchaux, M. Bourgoin, and A. Cartellier. Preferential concentration of heavy particles: A Voronoï analysis. *Phys. Fluids*, 22:103304, 2010.

[139] Robert D. Moser, John Kim, and Nagi N. Mansour. Direct numerical simulation of turbulent channel flow up to $re_\tau = 590$. *Physics of Fluids*, 11(4):943–945, 1999. doi: 10.1063/1.869966.

[140] Angelo Murrone and Hervé Guillard. A five equation reduced model for compressible two phase flow problems. *Journal of Computational Physics*, 202(2):664–698, 2005.

[141] R. Narasimha and K. R. Sreenivasan. Relaminarization in highly accelerated turbulent boundary layers. *Journal of Fluid Mechanics*, 61(3):417–447, 1973. doi: 10.1017/S0022112073000790.

[142] Ali Nassiri, Brad Kinsey, and Greg Chini. Shear instability of plastically-deforming metals in high-velocity impact welding. *Journal of the Mechanics and Physics of Solids*, 95:351–373, 2016. ISSN 0022-5096. doi: https://doi.org/10.1016/j.jmps.2016.06.002.

[143] S Ndanou, N Favrie, and S Gavrilyuk. Criterion of hyperbolicity in hyperelasticity in the case of the stored energy in separable form. *Journal of Elasticity*, 115(1):1–25, 2014.

[144] S Ndanou, Nicolas Favrie, and S Gavrilyuk. Multi-solid and multi-fluid diffuse interface model: Applications to dynamic fracture and fragmentation. *Journal of Computational Physics*, 295: 523–555, 2015.

[145] C. Nilsen, H. Andersson, and L. Zhao. A Voronoï analysis of preferential concentration in a vertical channel flow. *Phys. Fluids*, 25:115108, 2013.

[146] Robert R Nourgaliev, Truc-Nam Dinh, and Theo G Theofanous. Adaptive characteristics-based matching for compressible multifluid dynamics. *Journal of Computational Physics*, 213 (2):500–529, 2006.

[147] A López Ortega, M Lombardini, D I Pullin, and D I Meiron. Numerical simulations of the Richtmyer-Meshkov instability in solid-vacuum interfaces using calibrated plasticity laws. *Physical Review E*, 89(3):033018, 2014.

[148] A López Ortega, M Lombardini, DI Pullin, and Daniel I Meiron. Numerical simulation of elastic–plastic solid mechanics using an Eulerian stretch tensor approach and HLLD Riemann solver. *Journal of Computational Physics*, 257:414–441, 2014.

[149] A López Ortega, M Lombardini, P T Barton, D I Pullin, and D I Meiron. Richtmyer–Meshkov instability for elastic–plastic solids in converging geometries. *Journal of the Mechanics and Physics of Solids*, 76:291–324, 2015.

[150] Gustavo J. Otero R., Ashish Patel, Rafael Diez S., and Rene Pecnik. Turbulence modelling for flows with strong variations in thermo-physical properties. *International Journal of Heat and Fluid Flow*, 73:114–123, 2018. ISSN 0142-727X. doi: https://doi.org/10.1016/j.ijheatfluidflow. 2018.07.005.

[151] T. Oujia, K. Matsuda, and K. Schneider. Divergence and convergence of inertial particles in high-Reynolds-number turbulence. *J. Fluid Mech.*, 905:A14, 2020.

[152] T. Oujia, K. Matsuda, and K. Schneider. Extreme divergence and rotation values of the inertial particle velocity in high Reynolds number turbulence using Delaunay tessellation. In *12th International Symposium on Turbulence and Shear Flow Phenomena (TSFP12)*, Osaka, Japan, July 2022.

[153] T. Oujia, K. Matsuda, and K. Schneider. Computing differential operators of the particle velocity in moving particle clouds using tessellations. (in preparation), 2022. URL arXiv.2212.03580.

[154] Pedram Pakseresht and Sourabh V. Apte. Volumetric displacement effects in Euler-Lagrange LES of particle-laden jet flows. *International Journal of Multiphase Flow*, 113:16–32, 2019. ISSN 0301-9322. doi: https://doi.org/10.1016/j.ijmultiphaseflow.2018.12.013.

[155] Qingqing Pan, Hong Xiang, Ze Wang, Helge I. Andersson, and Lihao Zhao. Kinetic energy balance in turbulent particle-laden channel flow. *Physics of Fluids*, 32(7):073307, 2020. doi: 10.1063/5.0012570.

[156] Lorenzo Pareschi and Giovanni Russo. Implicit–explicit runge–kutta schemes and applications to hyperbolic systems with relaxation. *Journal of Scientific Computing*, 25(1):129–155, Oct 2005. ISSN 1573-7691. doi: 10.1007/s10915-004-4636-4.

[157] Marco Pasetto, Jonghyuk Baek, Jiun-Shyan Chen, Haoyan Wei, Jesse A. Sherburn, and Michael J. Roth. A lagrangian/semi-lagrangian coupling approach for accelerated meshfree modelling of extreme deformation problems. *Computer Methods in Applied Mechanics and Engineering*, 381:113827, 2021. ISSN 0045-7825. doi: https://doi.org/10.1016/j.cma.2021.113827.

[158] Ashish Patel, Bendiks J. Boersma, and Rene Pecnik. The influence of near-wall density and viscosity gradients on turbulence in channel flows. *Journal of Fluid Mechanics*, 809:793–820, 2016. doi: 10.1017/jfm.2016.689.

[159] Rene Pecnik and Ashish Patel. Scaling and modelling of turbulence in variable property channel flows. *Journal of Fluid Mechanics*, 823:R1, 2017. doi: 10.1017/jfm.2017.348.

[160] Guillaume Perigaud and Richard Saurel. A compressible flow model with capillary effects. *Journal of Computational Physics*, 209(1):139–178, 2005.

[161] Bradley J Plohr and David H Sharp. A conservative Eulerian formulation of the equations for elastic flow. *Advances in Applied Mathematics*, 9(4):481–499, 1988.

[162] Bradley J. Plohr and David H. Sharp. A conservative formulation for plasticity. *Advances in Applied Mathematics*, 13(4):462–493, 1992.

[163] Stephen B Pope. *Turbulent flows*. Cambridge university press, 2000.

[164] Hadi Pouransari and Ali Mani. Effects of Preferential Concentration on Heat Transfer in Particle-Based Solar Receivers. *Journal of Solar Energy Engineering*, 139(2), 11 2016. ISSN 0199-6231. doi: 10.1115/1.4035163. 021008.

[165] Hadi Pouransari and Ali Mani. Particle-to-fluid heat transfer in particle-laden turbulence. *Phys. Rev. Fluids*, 3:074304, Jul 2018. doi: 10.1103/PhysRevFluids.3.074304.

[166] M. Priestley. *Spectral Analysis and Time Series*. Academic, 1981.

[167] James J Quirk and Smadar Karni. On the dynamics of a shock–bubble interaction. *Journal of Fluid Mechanics*, 318:129–163, 1996.

[168] M. Reeks. The transport of discrete particles in inhomogeneous turbulence. *J. Aerosol Sci*, 14:729–739, 1983.

[169] David H. Richter and Peter P. Sullivan. Momentum transfer in a turbulent, particle-laden Couette flow. *Physics of Fluids*, 25(5):053304, 2013. doi: 10.1063/1.4804391.

[170] J. R. Rumble. *CRC handbook of Chemistry and Physics*. CRC Press, Boca Raton, FL, 98th edition, 2017.

[171] Grant Rydquist and Mahdi Esmaily. An optimal o(n) scheme for simulations of colliding, particle-laden flows on unstructured grids. *Journal of Computational Physics*, 420:109703, 2020. ISSN 0021-9991. doi: https://doi.org/10.1016/j.jcp.2020.109703.

[172] Lionel Sainsaulieu. Finite volume approximation of two phase-fluid flows based on an approximate Roe-type Riemann solver. *Journal of Computational Physics*, 121(1):1–28, 1995.

[173] G. Sardina, P. Schlatter, L. Brandt, F. Picano, and C. M. Casciola. Wall accumulation and spatial localization in particle-laden wall flows. *Journal of Fluid Mechanics*, 699:50–78, 2012. doi: 10.1017/jfm.2012.65.

[174] Richard Saurel and Rémi Abgrall. A simple method for compressible multifluid flows. *SIAM Journal on Scientific Computing*, 21(3):1115–1145, 1999.

[175] Richard Saurel and Carlos Pantano. Diffuse-interface capturing methods for compressible two-phase flows. *Annual Review of Fluid Mechanics*, 50(1), 2018.

[176] Richard Saurel, Fabien Petitpas, and Ray A Berry. Simple and efficient relaxation methods for interfaces separating compressible fluids, cavitating flows and shocks in multiphase mixtures. *Journal of Computational Physics*, 228(5):1678–1712, 2009.

[177] L. Schiller and Z. Naumann. A drag coefficient correlation. *Zeit. Ver. Deutsch Ing.*, 77(318): 318–320, 1935.

[178] Chi-Wang Shu and Stanley Osher. Efficient implementation of essentially non-oscillatory shock-capturing schemes, II. *Journal of Computational Physics*, 83(1):32–78, 1989. ISSN 0021-9991. doi: https://doi.org/10.1016/0021-9991(89)90222-2.

[179] Ratnesh K Shukla. Nonlinear preconditioning for efficient and accurate interface capturing in simulation of multicomponent compressible flows. *Journal of Computational Physics*, 276: 508–540, 2014.

[180] Ratnesh K Shukla, Carlos Pantano, and Jonathan B Freund. An interface capturing method for the simulation of multi-phase compressible flows. *Journal of Computational Physics*, 229 (19):7411–7439, 2010.

[181] Keh-Ming Shyue. An efficient shock-capturing algorithm for compressible multicomponent problems. *Journal of Computational Physics*, 142(1):208–242, 1998.

[182] Keh-Ming Shyue and Feng Xiao. An Eulerian interface sharpening algorithm for compressible two-phase flow: The algebraic THINC approach. *Journal of Computational Physics*, 268: 326–354, 2014.

[183] Nathan P Siegel, Clifford K Ho, Siri S Khalsa, and Gregory J Kolb. Development and evaluation of a prototype solid particle receiver: on-sun testing and model validation. *Journal of Solar Energy Engineering*, 132(2), 2010.

[184] KK So, XY Hu, and NA Adams. Anti-diffusion interface sharpening technique for two-phase compressible flow simulations. *Journal of Computational Physics*, 231(11):4304–4323, 2012.

[185] K. Squires and J. Eaton. Preferential concentration of particles by turbulence. *Phys. Fluids A*, 3:1169–1178, 1991.

[186] DJ Steinberg, SG Cochran, and MW Guinan. A constitutive model for metals applicable at high-strain rate. *Journal of Applied Physics*, 51(3):1498–1504, 1980.

[187] Akshay Subramaniam. *Simulations of Shock Induced Interfacial Instabilities Including Materials with Strength*. PhD thesis, Stanford University, 2018.

[188] Akshay Subramaniam, Niranjan S Ghaisas, and Sanjiva K Lele. High-order Eulerian simulations of multimaterial elastic–plastic flow. *Journal of Fluids Engineering*, 140(5):050904, 2018.

[189] Kazuyasu Sugiyama, Satoshi Ii, Shintaro Takeuchi, Shu Takagi, and Yoichiro Matsumoto. Full Eulerian simulations of biconcave neo-Hookean particles in a Poiseuille flow. *Computational Mechanics*, 46(1):147–157, 2010.

[190] Kazuyasu Sugiyama, Satoshi Ii, Shintaro Takeuchi, Shu Takagi, and Yoichiro Matsumoto. A full Eulerian finite difference approach for solving fluid–structure coupling problems. *Journal of Computational Physics*, 230(3):596–627, 2011.

[191] William Sutherland. LII. the viscosity of gases and molecular force. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 36(223):507–531, 1893. doi: 10.1080/14786449308620508.

[192] Geoffrey Ingram Taylor. The use of flat-ended projectiles for determining dynamic yield stress I. theoretical considerations. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 194(1038):289–299, 1948. doi: 10.1098/rspa.1948.0081.

[193] Hendrik Tennekes and John Leask Lumley. *A First Course in Turbulence*. MIT Press, 1972.

[194] Hiroshi Terashima and Grétar Tryggvason. A front-tracking/ghost-fluid method for fluid interfaces in compressible flows. *Journal of Computational Physics*, 228(11):4012–4037, 2009.

[195] VA Titarev, E Romenski, and EF Toro. Musta-type upwind fluxes for non-linear elasticity. *International Journal for Numerical Methods in Engineering*, 73(7):897–926, 2008.

[196] Arpit Tiwari, Jonathan B Freund, and Carlos Pantano. A diffuse interface model with immiscibility preservation. *Journal of Computational Physics*, 252:290–309, 2013.

[197] John A Trangenstein and Phillip Colella. A higher-order Godunov method for modeling finite deformation in elastic–plastic solids. *Communications on Pure and Applied Mathematics*, 44(1):41–100, 1991.

[198] Andrew Trettel and Johan Larsson. Mean velocity scaling for compressible wall turbulence with heat transfer. *Physics of Fluids*, 28(2):026102, 2016. doi: 10.1063/1.4942022.

[199] V K Tritschler, B J Olson, S K Lele, Stefan Hickel, X Y Hu, and Nikolaus Andreas Adams. On the Richtmyer–Meshkov instability evolving from a deterministic multimode planar interface. *Journal of Fluid Mechanics*, 755:429–462, 2014.

[200] Boris Valkov, Chris H. Rycroft, and Ken Kamrin. Eulerian Method for Multiphase Interactions of Soft Solid Bodies in Fluids. *Journal of Applied Mechanics*, 82(4), 04 2015. ISSN 0021-8936. doi: 10.1115/1.4029765. 041011.

[201] E. R. Van Driest. Turbulent boundary layer in compressible fluids. *Journal of the Aeronautical Sciences*, 18(3):145–160, 1951. doi: 10.2514/8.1895.

[202] Sankaran Venkateswaran, Jules W Lindau, Robert F Kunz, and Charles L Merkle. Computation of multiphase mixture flows with compressibility effects. *Journal of Computational Physics*, 180(1):54–77, 2002.

[203] A. W. Vreman. Turbulence characteristics of particle-laden pipe flow. *Journal of Fluid Mechanics*, 584:235–279, 2007. doi: 10.1017/S0022112007006556.

[204] Bert Vreman, Bernard J Geurts, NG Deen, JAM Kuipers, and Johannes GM Kuerten. Two- and four-way coupled Euler–Lagrangian large-eddy simulation of turbulent particle-laden channel flow. *Flow, Turbulence and Combustion*, 82(1):47–71, 2009.

[205] J. R. West, T. Oujia, K. Matsuda, K. Schneider, S. S. Jain, and K. Maeda. Divergence and curl of the inertial particle velocity in a four-way coupled turbulent channel flow. In *Proceedings of the Summer Program, Center for Turbulence Research, Stanford University*, 2022.

[206] J. R. West, M. C. Adler, and S. K. Lele. A high-order, localized artificial diffusivity method for Eulerian simulation of multi-material elastic-plastic deformation with strain hardening. (in preparation), 2023.

[207] J. R. West, T. Oujia, K. Matsuda, K. Schneider, S. S. Jain, and K. Maeda. Clustering, rotation, and swirl of inertial particles in turbulent channel flow. Manuscript submitted for publication to the Journal of Fluid Mechanics, 2023.

[208] Mark L. Wilkins and Michael W. Guinan. Impact of cylinders on a rigid boundary. *Journal of Applied Physics*, 44(3):1200–1206, 1973. doi: 10.1063/1.1662328.

[209] M Wilkinson and B Mehlig. Caustics in turbulent aerosols. *Europhysics Letters (EPL)*, 71(2): 186–192, July 2005. doi: 10.1209/epl/i2004-10532-7.

[210] Man Long Wong and Sanjiva K Lele. High-order localized dissipation weighted compact nonlinear scheme for shock-and interface-capturing in compressible flows. *Journal of Computational Physics*, 339:179–209, 2017.

[211] Yan Xia, Zhaowu Lin, Dingyi Pan, and Zhaosheng Yu. Turbulence modulation by finite-size heavy particles in a downward turbulent channel flow. *Physics of Fluids*, 33(6):063321, 2021. doi: 10.1063/5.0053540.

[212] Y. Yamamoto, M. Potthoff, T. Tanaka, T. Kajishima, and Y. Tsuji. Large-eddy simulation of turbulent gas–particle flow in a vertical channel: effect of considering inter-particle collisions. *Journal of Fluid Mechanics*, 442:303–334, 2001. doi: 10.1017/S0022112001005092.

[213] Geum-Su Yeom and Keun-Shik Chang. A modified HLLC-type Riemann solver for the compressible six-equation two-fluid model. *Computers & Fluids*, 76:86–104, 2013.

[214] Zhaosheng Yu, Yan Xia, Yu Guo, and Jianzhong Lin. Modulation of turbulence intensity by heavy finite-size particles in upward channel flow. *Journal of Fluid Mechanics*, 913:A3, 2021. doi: 10.1017/jfm.2020.1140.

[215] R. Zamansky, F. Coletti, M. Massot, and A. Mani. Turbulent thermal convection driven by heated inertial particles. *Journal of Fluid Mechanics*, 809:390–437, 2016. doi: 10.1017/jfm. 2016.630.

[216] You-Sheng Zhang, Wei-Tao Bi, Fazle Hussain, Xin-Liang Li, and Zhen-Su She. Mach-number-invariant mean-velocity profile of compressible turbulent boundary layers. *Phys. Rev. Lett.*, 109:054502, July 2012. doi: 10.1103/PhysRevLett.109.054502.

[217] Yuan Zhang, Sudarsanam Suresh Babu, Curtis Prothe, Michael Blakely, John Kwasegroch, Mike LaHa, and Glenn S. Daehn. Application of high velocity impact welding at varied different length scales. *Journal of Materials Processing Technology*, 211(5):944–952, 2011. ISSN 0924-0136. doi: https://doi.org/10.1016/j.jmatprotec.2010.01.001. Special Issue: Impulse Forming.

[218] Lihao Zhao, Helge I. Andersson, and Jurriaan J. J. Gillissen. Interphasial energy transfer and particle dissipation in particle-laden wall turbulence. *Journal of Fluid Mechanics*, 715:32–59, 2013. doi: 10.1017/jfm.2012.492.

[219] Tian Zhou, Lihao Zhao, Weixi Huang, and Chunxiao Xu. Non-monotonic effect of mass loading on turbulence modulations in particle-laden channel flow. *Physics of Fluids*, 32(4): 043304, 2020. doi: 10.1063/5.0002114.

[220] Ye Zhou. Rayleigh–Taylor and Richtmyer–Meshkov instability induced flow, turbulence, and mixing. I. *Physics Reports*, 720:1–136, 2017.