Numerical methods for self-consistent dust-gas dynamics in hydrodynamical simulations

Giovanni Tedeschi Prades

München 2026

Numerical methods for self-consistent dust-gas dynamics in hydrodynamical simulations

Giovanni Tedeschi Prades

Dissertation
der Fakultät für Physik
der Ludwig-Maximilians-Universität
München

vorgelegt von
Giovanni Tedeschi Prades
aus Rom, Italien

München, den 18. Dezember 2025

Erstgutachter: Prof. Dr. Til Birnstiel
Zweitgutachter: Dr. Klaus Dolag
Tag der mündlichen Prüfung: 12. Februar 2026

"Knowledge is night!"

Professor Nightingale, in *The 13½ Lives of Captain Bluebear*, Walter Moers

The work contained in this thesis was carried out at the University Observatory of the Ludwig Maximilian University of Munich, under the supervision of Prof. Dr. Til Birnstiel, Dr. Klaus Dolag, and Prof. Dr. Barbara Ercolano. The committee members of the doctoral defence were Prof. Dr. Til Birnstiel, Dr. Klaus Dolag, Prof. Dr. Volker Springel, and Prof. Dr. Barbara Ercolano.
I acknowledge the use of large language model–based tools for assistance with grammar, wording, and stylistic refinement of the text.

Contents

List of Publications

First author publications

- **Modeling dust dynamics in `OpenGadget3` - I. SPH implementation of the One-Fluid model**
  Giovanni Tedeschi-Prades, Til Birnstiel; Klaus Dolag, Barbara Ercolano, and Mark Hutchison.
  Astronomy & Astrophysics, Volume 704, n. A92 (2025)

- **General Implicit Runge-Kutta Integrators for Multifluid Gas-Dust Aerodynamic Drag**
  Giovanni Tedeschi-Prades, Til Birnstiel; Klaus Dolag, Barbara Ercolano, and Mark Hutchison.
  The Astrophysical Journal, Volume 994, n. 267 (2025)

- **Modeling dust dynamics in `OpenGadget3` - II. A Meshless Finite Mass implementation of the multi-grain One–Fluid model**
  Giovanni Tedeschi-Prades, Til Birnstiel; Klaus Dolag, Barbara Ercolano, Mark Hutchison, and Frederick Groth.
  in prep. for submission on Astronomy & Astrophysics

Co-authored publications

- **Full one-fluid dusty gas with multiple grain species in SPH**
  Mark Hutchison, Guillaume Laibe, **Giovanni Tedeschi-Prades**, Timothée David-Cléris, Alexandre Barret, Maxime Lombart, Daniel J. Price, and Christine M. Koepferl.
  submitted to Monthly Notices of the Royal Astronomical Society

- **OpenGadget3**
  Klaus Dolag, Geray S. Karademir, Luca Tornatore, Antonio Ragagnin, Milena Valentini, Ludwig Böss, Frederick Groth, ..., **Giovanni Tedeschi-Prades**.
  in prep. for submission on Astronomy and Computing

List of Figures

List of Tables

Abstract

Cosmic dust is a fundamental constituent of astrophysical systems, playing a key role in the thermal balance, chemistry, and observable properties of the interstellar medium, star–forming regions, protoplanetary disks, and galaxies. While often treated as a passive tracer perfectly coupled to the gas, dust can in many environments exhibit dynamical behavior that differs significantly from that of the gas, with important consequences for both observations and theoretical models.

The One–Fluid model provides an efficient framework to describe dust–gas mixtures by evolving a single set of fluid variables, while allowing for relative motion between the phases. However, most existing implementations of this model rely on the terminal velocity approximation, which restricts its applicability to tightly coupled dust grains and prevents the exploration of regimes characterized by weak dust–gas coupling or high dust–to–gas ratios.

In this thesis, I pushed the boundaries of simulating dust dynamics in hydrodynamical simulations beyond the terminal velocity approximation. I achieved this by improving existing numerical formulations and by developing new methods compatible with modern hydrodynamical simulation frameworks. My work spans both Smoothed Particle Hydrodynamics (SPH) and Meshless Finite Mass (MFM) methods.

In the SPH framework, I improved upon the original One–Fluid model by deriving and implementing a discretization of the full system of equations, lifting the terminal velocity approximation. This enabled the self–consistent simulation of loosely coupled dust fluids within the One–Fluid formalism. I also identified a fundamental limitation shared by all fluid–based dust models, namely, the absence of interpenetrating dust trajectories, and introduced a diffusion–based regularization scheme to stabilize the dynamics in regimes where this limitation becomes numerically problematic.

To efficiently integrate aerodynamic drag for multiple dust species, I developed a new implicit drag integrator, the General Implicit Runge–Kutta (GIRK) method. GIRK is compatible with arbitrary Strang–splitting schemes, converges to the correct equilibrium solution with second–order accuracy, and scales linearly with the number of dust species, making it well suited for simulations with complex grain size distributions.

Building upon these developments, I implemented a self–consistent treatment of dust dynamics for multiple dust species in the MFM framework of `OpenGadget3`. In this approach, simulation particles represent the dust–gas mixture as in the One–Fluid model, while gas and dust are treated separately at the level of the Riemann solver. To satisfy the strict mass–conservation requirements of MFM, I developed a novel algorithm to determine the zero–mass–flux reference frame for multi–fluid interactions. This hybrid formulation combines the advantages of the One–Fluid model with the accuracy of Godunov–type solvers and avoids the need for artificial viscosity terms.

The methods developed in this thesis are validated against a suite of standard benchmark tests and provide a flexible and robust framework for simulating dust dynamics across a wide range of coupling regimes. Moreover, I conducted SPH simulations of complex flows,

such as the Kelvin–Helmholtz instability, to showcase the distinct behaviors of gas and dust and the impact of large dust–to–gas ratios on gas dynamics. These methods open new possibilities for studying the role of dust in astrophysical systems, from interstellar and galactic environments to star and planet formation, and offer a solid foundation for future extensions, including magnetohydrodynamics, dust growth and fragmentation, and more advanced models of dust dynamics.

Zusammenfassung

Kosmischer Staub ist ein grundlegender Bestandteil astrophysikalischer Systeme und spielt eine zentrale Rolle für das thermische Gleichgewicht, die Chemie und die beobachtbaren Eigenschaften des interstellaren Mediums, von Sternentstehungsgebieten, protoplanetaren Scheiben und Galaxien. Obwohl Staub häufig als passiver Tracer behandelt wird, der perfekt an das Gas gekoppelt ist, kann er in vielen Umgebungen ein dynamisches Verhalten zeigen, das sich deutlich von dem des Gases unterscheidet, mit wichtigen Konsequenzen sowohl für Beobachtungen als auch für theoretische Modelle.
Das Ein–Fluid–Modell bietet einen effizienten Modellierungsansatz zur Beschreibung von Staub–Gas–Gemischen, indem ein einziger Satz von Fluidvariablen entwickelt wird, während relative Bewegungen zwischen den Phasen zugelassen werden. Die meisten existierenden Implementierungen dieses Modells beruhen jedoch auf der Terminalgeschwindigkeitsnäherung, die seine Anwendbarkeit auf eng gekoppelte Staubkörner beschränkt und die Untersuchung von Parameterbereichen mit schwacher Staub–Gas–Kopplung oder hohen Staub–zu–Gas–Verhältnissen verhindert.
In dieser Dissertation habe ich die Grenzen der Simulation von Staubdynamik in hydrodynamischen Simulationen über die Terminalgeschwindigkeitsnäherung hinaus verschoben. Dies gelang mir durch die Verbesserung bestehender numerischer Formulierungen sowie durch die Entwicklung neuer Methoden, die mit modernen hydrodynamischen Simulationsmethoden kompatibel sind. Meine Arbeit umfasst sowohl Smoothed Particle Hydrodynamics (SPH) als auch Meshless Finite Mass (MFM)–Methoden. Im SPH–Rahmen habe ich das ursprüngliche Ein–Fluid–Modell verbessert, indem ich eine Diskretisierung des vollständigen Gleichungssystems hergeleitet und implementiert habe, wodurch die Terminalgeschwindigkeitsnäherung aufgehoben wurde. Dies ermöglichte die selbstkonsistente Simulation locker gekoppelter Staubfluide innerhalb des Ein–Fluid–Formalismus. Zudem identifizierte ich eine grundlegende Einschränkung, die allen fluidbasierten Staubmodellen gemeinsam ist, nämlich das Fehlen sich durchdringender Staubtrajektorien, und führte ein diffusionsbasiertes Regularisierungsschema ein, um die Dynamik in Parameterbereichen zu stabilisieren, in denen diese Einschränkung numerisch problematisch wird.
Zur effizienten Integration des aerodynamischen Widerstands für mehrere Staubarten entwickelte ich einen neuen impliziten Integrator für den Widerstandsterm, die General Implicit Runge–Kutta (GIRK)–Methode. GIRK ist mit beliebigen Strang–Splitting–Schemata kompatibel, konvergiert mit Genauigkeit zweiter Ordnung zur korrekten Gleichgewichtslösung und skaliert linear mit der Anzahl der Staubarten, was sie besonders geeignet für Simulationen mit komplexen Korngrößenverteilungen macht.
Aufbauend auf diesen Entwicklungen implementierte ich eine selbstkonsistente Behandlung der Staubdynamik für mehrere Staubarten im MFM–Rahmen von `OpenGadget3`. In diesem Ansatz repräsentieren Simulationspartikel das Staub–Gas–Gemisch wie im Ein–Fluid–Modell, während Gas und Staub auf der Ebene des Riemann–Lösers getrennt behandelt werden. Um die strengen Anforderungen der MFM–Methode an die Massenerhaltung zu erfüllen, entwickelte ich einen neuartigen Algorithmus zur Bestimmung des Bezugsrahmens

ohne Massenfluss für Multi–Fluid–Interaktionen. Diese hybride Formulierung vereint die Vorteile des Ein–Fluid–Modells mit der Genauigkeit von Godunov–artigen Lösern und vermeidet die Notwendigkeit künstlicher Viskositätsterme.
Die in dieser Dissertation entwickelten Methoden werden anhand einer Reihe standardisierter Benchmark–Tests validiert und bieten einen flexiblen und robusten Rahmen für die Simulation der Staubdynamik über einen weiten Bereich der Kopplung hinweg. Darüber hinaus habe ich SPH–Simulationen komplexer Strömungen, wie etwa der Kelvin–Helmholtz–Instabilität, durchgeführt, um die unterschiedlichen Verhaltensweisen von Gas und Staub sowie den Einfluss hoher Staub–zu–Gas–Verhältnisse auf die Gasdynamik aufzuzeigen. Diese Methoden eröffnen neue Möglichkeiten zur Untersuchung der Rolle von Staub in astrophysikalischen Systemen, vom interstellaren und galaktischen Umfeld bis hin zur Stern– und Planetenentstehung, und bieten eine solide Grundlage für zukünftige Erweiterungen, einschließlich Magnetohydrodynamik, Staubwachstum und –fragmentierung sowie weiterentwickelter Modelle der Staubdynamik.

1

Introduction

Cosmic dust is a ubiquitous yet complex component of astrophysical systems, playing a central role in processes that span an extraordinary range of physical scales, from the microphysics of solid-state interactions to the macroscopic evolution of galaxies, stars, and planetary systems. Although dust typically accounts for only about one percent of the mass of the interstellar medium (ISM), its presence fundamentally alters the thermal, chemical, and dynamical evolution of gas. By efficiently absorbing, scattering, and re-emitting radiation, dust regulates the energy balance of astrophysical environments, enables the formation of molecules on its surface, and provides the raw material from which planets ultimately form.

Dust grains originate primarily in the outflows of evolved stars and in supernova ejecta, where refractory elements condense into solid particles under extreme physical conditions. Once injected into the ISM, however, dust does not remain static. Instead, it embarks on a complex evolutionary journey shaped by a wide range of physical processes, including fragmentation in interstellar shocks, coagulation in dense regions, accretion of gas-phase metals, and partial destruction by energetic radiation and collisions. As a result, the population of dust grains observed in the diffuse ISM differs markedly from the dust produced at stellar sources, both in composition and in size distribution. Understanding this evolution is essential for interpreting observations of extinction, emission, and polarization across the electromagnetic spectrum.

Beyond its radiative impact, dust plays a decisive role in the physics of star formation. In molecular clouds, dust grains act as an efficient coolant by re-radiating thermal energy in the infrared (IR), thereby lowering the gas temperature and facilitating gravitational collapse. Dust surfaces also catalyze the formation of molecular hydrogen, the dominant constituent of cold star-forming gas, linking solid-state physics directly to the onset of star formation. The abundance, size distribution, and composition of dust therefore influence the characteristic mass scale of star formation and the fragmentation of molecular clouds.

As stars form and accrete circumstellar material, dust continues to shape the evolution of their surrounding environments. In protoplanetary disks, dust grains interact aerodynamically with gas, migrate radially, settle toward the disk midplane, and grow through successive collisions. These processes determine both the spatial distribution of solids and the maximum attainable grain sizes before growth barriers, such as radial drift and fragmentation, intervene. Since dust grains constitute the building blocks of planetesimals and, ultimately, planets, their dynamical and collisional evolution sets the initial conditions for planet formation. Observations of disks at IR and millimeter wavelengths, particularly with facilities such as the Atacama Large Millimeter Array (ALMA), have revealed striking substructures that are widely interpreted as signatures of dust evolution and transport.
The first section of this chapter introduces the main observational techniques used to infer the physical and chemical properties of interstellar dust, such as its abundance, composition, and grain size distribution (GSD). By combining extinction, emission, and scattering measurements across a broad range of wavelengths, these tools provide indirect but powerful constraints on the nature of dust in different astrophysical environments.
The second section focuses on the life cycle of dust, following its evolution from its formation in stellar sources to its processing in the diffuse ISM, and finally to its role during star and planet formation. Particular emphasis is placed on how dust properties change across these environments, as grains undergo growth, destruction, and dynamical evolution. This evolutionary perspective highlights the central role of dust as a link between stellar evolution, the physics of the ISM, and the early stages of planetary system formation.

1.1 Observational evidences of dust properties

This section briefly reviews several observational tools used to infer the presence of interstellar dust and to study its properties. I focus in particular on extinction curves, IR emission, and the analysis of presolar grains. This list is not exhaustive; other valuable methods, such as scattered-light observations and polarization studies, are not discussed here.
For illustrative purposes, each observable is associated with one example of a dust property that can be inferred from it. These associations are not unique but serve only as representative cases. Moreover, the discussion is limited to a subset of dust properties: grain-size distribution, total dust mass, and chemical composition. In reality, dust characteristics are far more complex, involving grain shape, charge, porosity, temperature, and many other parameters.
A further complication arises from the fact that dust exists in a wide range of environments that differ in density, pressure, temperature, and ambient radiation field, from the diffuse ISM of the Milky Way to dense molecular clouds, and from high-redshift galaxies to quiescent elliptical systems. For these reasons, the summary presented here necessarily restricts itself to selected observables and properties, with particular emphasis on the diffuse ISM.

1.1.1 Extinction curves and GSD

The existence of interstellar dust was first inferred by Trumpler (1930) to explain the phenomenon of stellar reddening. When comparing the spectra of two stars of the same spectral class and luminosity along different lines of sight, one would expect to obtain very similar spectral shapes. However, in practice, one spectrum often appears more reddened than the other, with less flux at shorter wavelengths. Among the explanations suggested by Trumpler (1930) was the presence of "fine cosmic dust particles of various sizes maintained in space by the light pressure of the stars, producing the observed selective absorption by Rayleigh scattering."
The wavelength-dependent extinction A_λ can be quantified by comparing the observed flux from a given star (F_λ), to the flux of a star of the same spectral type but in the absence of extinction (F_λ^0):

$$A_\lambda = 2.5 \log_{10}(F_\lambda^0/F_\lambda). \tag{1.1}$$

The pair method, however, has several drawbacks. To name two, it is sensitive to even tiny differences in the spectral properties of the two stars used, and it is impractical for rare stars, such as O-type stars or supergiants. For this reason, more modern studies of extinction curves rely on the method by Fitzpatrick & Massa (2007), in which the reference flux F_λ^0 is replaced by a modelled flux F_λ^{model}, obtained by fitting the spectral energy distribution (SED) of the observed star.
Fig. 1.1 shows the average normalized extinction curve for the Milky Way, obtained by Fitzpatrick & Massa (2007), averaging 328 extinction curves along different lines of sight from different stellar sources.

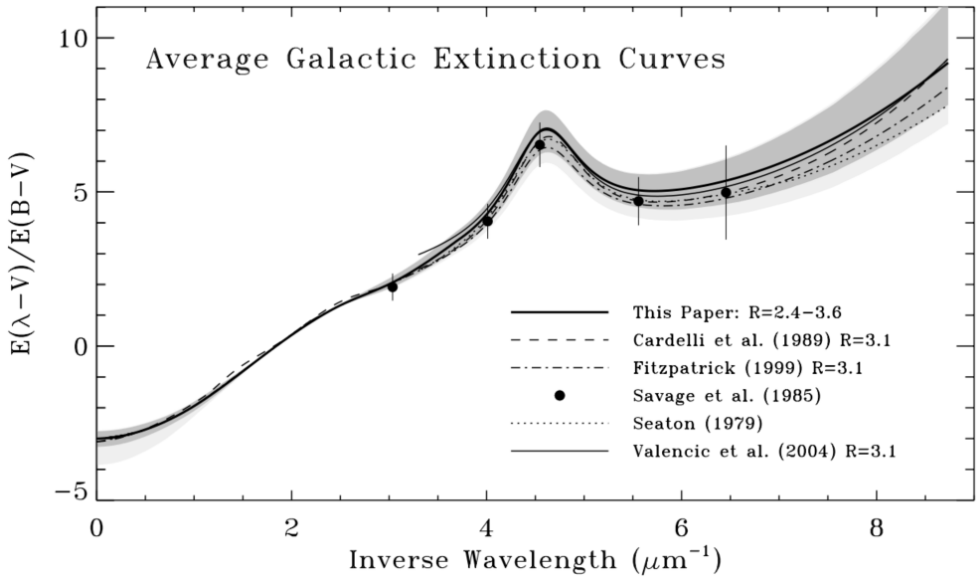

Figure 1.1: Average galactic extinction curve from Fitzpatrick & Massa (2007) for the Milky Way as a function of the inverse wavelength.

The first feature in the typical Milky Way extinction curve is the "2175 Å bump" (4.6 μm^{-1}), generally attributed to carbonaceous grains such as polycyclic aromatic hydrocarbons (PAHs). This feature is evident in the extinction curves of the Milky Way, due to carbonaceous grains that survive long enough to form PAHs. In a different environemnt such as the Small Magellanic Clouds, which has lower metallicty ($Z \simeq 0.2 Z_\odot$) and higher interstellar UV field, PAHs are destroyed much more efficiently, and as a consequence their extinction curves do not show such a bump (Weingartner & Draine, 2001; Nozawa & Fukugita, 2013).
The second feature of the extinction curve is its slope, R_V, which can be defined as

$$R_V = \frac{A_V}{A_B - A_V}. \tag{1.2}$$

Interstellar dust is mostly transparent to IR radiation, while it absorbs most of the ultraviolet (UV) photons. Large grains, which absorb and scatter radiation over a broad wavelength range, produce flatter extinction curves with high R_V values, while a distribution dominated by small grains yields a steep UV rise, approaching the Rayleigh scattering limit $A_\lambda \propto \lambda^{-4}$ and lower R_V values. Zhang et al. (2023) found, using a sample of 3 million stars, a value of $R_V \simeq 3.24 \pm 0.34$. Given their large sample, they were also able to produce a detailed map of R_V for the galactic disk, shown in Fig. 1.2.

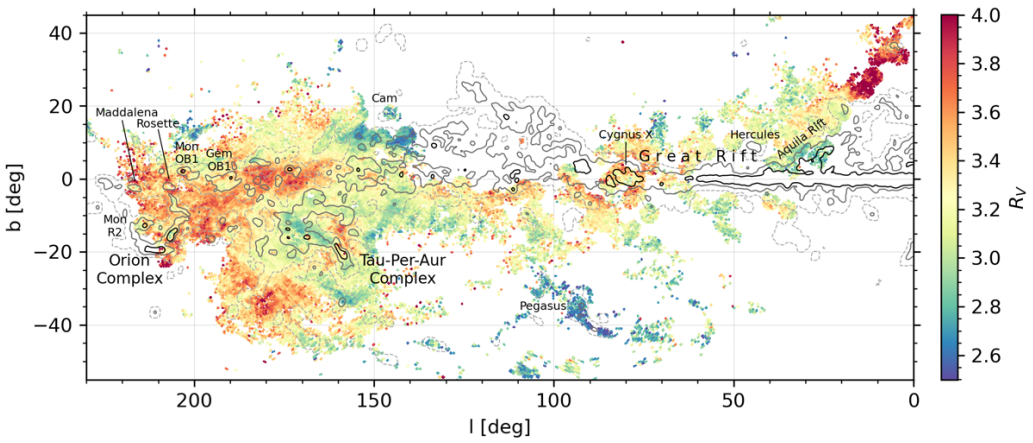

Figure 1.2: Spatial distribution of the extinction curve slope R_V on the galactic plane, from Zhang et al. (2023) obtained from the LAMOST telescope. The grey contour lines show the smooth CO map from Planck Collaboration et al. (2016), with darker lines corresponding to higher CO values. White areas are either unobserved or non-reliable regions.

This map shows that molecular clouds exhibit systematically lower R_V values than their surroundings. This suggests that, in molecular clouds, dust grains differ in their chemical and size properties compared to those in the diffuse ISM.
The extinction curves can be modelled and fitted using theoretical models, which are highly sensitive to the chosen GSD. The GSD, represented as $n(a)$, is defined such that $n(a)da$

is the number of particles with radius between a and $a + da$. Given the extinction law for each particle size $Q_{\rm ext}(a, \lambda)$, we can write the total extinction as an integral over all grain sizes

$$A_\lambda \propto \int_0^\infty n(a)\pi a^2 Q_{\rm ext}(a, \lambda) da. \tag{1.3}$$

Using such a parametrization, Mathis et al. (1977) where able to constrain the GSD in the diffuse ISM to be a power law

$$n(a) \propto a^{-q}, \tag{1.4}$$

with $q \approx 3.5$, for particles with radius 0.005 μm $< a <$ 0.25 μm.
From this size distribution, we can learn a few things about the properties of dust in the diffuse ISM. By integrating this distribution, we can obtain the fraction of the total population having a size smaller than a

$$f(a) = \frac{\int_{a_{\rm min}}^{a} n(a)da}{\int_{a_{\rm min}}^{a_{\rm max}} n(a)da} = \frac{a^{-2.5} - a_{\rm min}^{-2.5}}{a_{\rm max}^{-2.5} - a_{\rm min}^{-2.5}}. \tag{1.5}$$

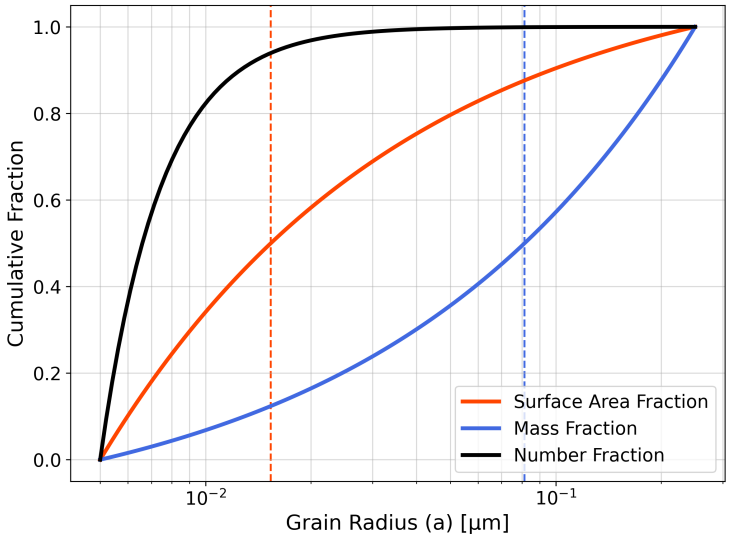

Figure 1.3: Cumulative fractions for the number of dust grains (in black), the surface area (in red), and the mass fraction (in blue), assuming an MRN GSD. Vertical dashed lines correspond to the 50% cumulative fractions of surface area (in red) and mass (in blue).

This cumulative function is plotted in Fig. 1.3 in black. The same exercise can be done for the cumulative mass and surface area as well, also plotted in Fig. 1.3. While the number of particles is dominated by the smaller population, the opposite is true for the total mass. Despite being less than 0.1% of the total number, the dust grains with size larger than 0.08 μm contain half of the total dust mass. As for the surface area, it is most skewed towards smaller grains, with half the total surface area in grains smaller than 0.015 μm, which account for roughly 95% of the total number of grains. These properties are useful

to keep in mind, since some processes, like gas accretion and surface chemistry, depend on the available surface area, while others, like thermal emission, depend on the total dust mass.

1.1.2 IR emission and dust mass

The probability that a photon of a given wavelength is absorbed by a dust grain of a specific composition and size is characterized by the absorption cross section $C_{\rm abs}(a, \lambda)$. After absorbing energy, the grain re-emits radiation approximately according to a Planck function $B_\lambda(T)$. However, a dust grain does not necessarily maintain a single temperature: small grains undergo stochastic heating and are better described by a temperature probability distribution dP/dT. By combining these processes, Draine & Li (2007) models the emissivity of a single dust grain as

$$j_\lambda = 4\pi \int C_{\rm abs}(\lambda) B_\lambda(T) \frac{dP}{dT} dT. \tag{1.6}$$

For large grains, which remain in thermal equilibrium with the local radiation field at a steady-state temperature $T_{\rm ss}$, the temperature distribution becomes sharply peaked, and can be approximated as

$$\frac{dP}{dT} \approx \delta(T - T_{\rm ss}). \tag{1.7}$$

The steady-state temperature is set by balancing the power absorbed from the ambient radiation field with the power radiated away by the grain.
Small grains, on the other hand, do not achieve thermal equilibrium. Even a single UV photon may raise their temperature by hundreds or thousands of Kelvin. Their temperature spikes and cools rapidly in a stochastic cycle that must be modeled statistically. A treatment of this process is given in Draine & Li (2001). The absorption cross section is related to the dust opacity κ_λ through

$$\kappa_\lambda = \frac{C_{\rm abs}(\lambda)}{m_{\rm grain}}, \tag{1.8}$$

where $m_{\rm grain}$ is the mass of an individual grain. Dust opacity depends on composition, size distribution, shape, porosity, and even temperature. Consequently, it is one of the dominant uncertainties in modeling IR dust emission. For illustrative purposes, we can use limiting results from Mie theory. Defining the size parameter

$$x = \frac{2\pi a}{\lambda}, \tag{1.9}$$

we find that for $x > 1$, or grains larger than the incoming wavelength, the absorption efficiency approaches a constant value, so that $C_{\rm abs}(\lambda)$ becomes roughly wavelength-independent. For $x < 1$, the Rayleigh limit, the absorption cross section scales approximately as $C_{\rm abs}(\lambda) \propto \frac{1}{\lambda}$. A detailed discussion of these limits is provided in Bohren &

Huffman (1998a,b).
Because of this, large grains emit similarly to blackbodies, while small grains exhibit emissivities weighted toward shorter wavelengths.
From the emissivity of the single grain, we can obtain the total emissivity of a cloud of dust

$$J_\lambda = N_\mathrm{d} \cdot j_\lambda, \tag{1.10}$$

where N_d is the number of dust grains. This linear relation implicitly assumes that each grain radiates independently, without multiple absorption or scattering of emitted photons. In other words, this requires an optically thin medium. If, in addition, all grains are approximated as being at a single temperature, we obtain

$$J_\lambda = 4\pi B_\lambda(T) \sum_i C_\mathrm{abs}(a_i, \lambda) = 4\pi B_\lambda(T) \kappa_\lambda M_\mathrm{d}, \tag{1.11}$$

where in the last step we used the definition of opacity in Eq. 1.8.
The observed flux is then

$$F_\lambda = \kappa_\lambda \cdot B_\lambda(T_\mathrm{d}) \cdot M_\mathrm{d} \cdot \frac{1}{D^2}, \tag{1.12}$$

where D is the distance between the observer and the dust cloud.
By inverting this expression, Hildebrand (1983) used IR flux measurements to estimate dust masses. Comparison with independent gas-mass estimates yields a Milky Way dust-to-gas ratio of roughly 0.01.

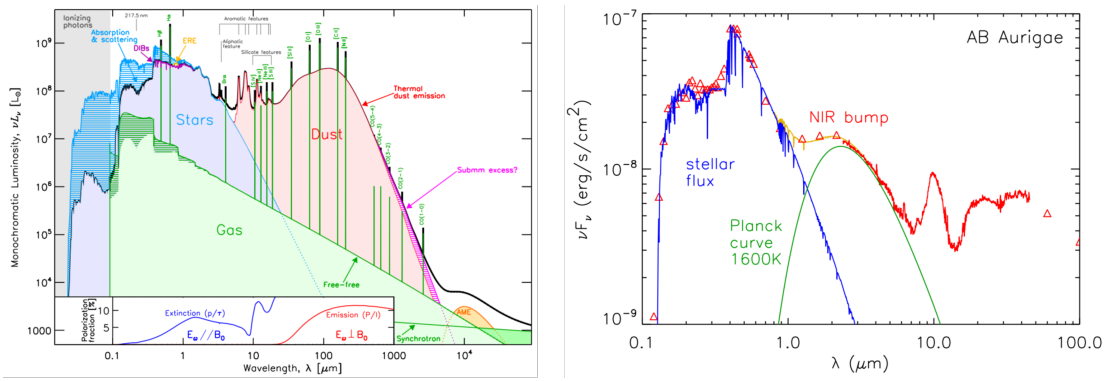

Figure 1.4: Spectral energy distributions of an entire late-type spiral galaxy (left), from Galliano et al. (2018), and the AB Aurigae star (right), from Dullemond & Monnier (2010).

Although conceptually simple, this approach is highly sensitive to assumptions about dust temperature and opacity. A more comprehensive method is to fit the full spectral energy distribution (SED). Fig. 1.4 shows examples of SEDs for a typical late-type galaxy (Galliano et al. (2018)) and for the Herbig Ae star AB Aurigae (Dullemond & Monnier (2010)). For an entire galaxy, the stellar part of the SED contains the contributions of stars of all temperatures. The dust, being at an equilibrium temperature of around $15 - 25\ K$, emits as a blackbody peaking around 100 μm. The thermal emission from dust is the combination of the emission of multiple dust grain sizes, each with its own emissivity modelled by

Eq. 1.6. Dust contributes to shaping the SED of a galaxy also via extinction, both via absorption and scattering. The role of extinction is shown in Fig. 1.4 as the blue hatched area. Dust emission in the IR also exhibits spectral features in addition to the continuum black-body component. In particular, PAHs contribute to the spectral features between $\sim 2 - 17$ μm, while amorphous silicates show a distinct peak at 10 μm. The SED of AB Aurigae, in panel b of Fig. 1.4 shows both observations (in red) and models of the stellar flux (in blue) and of dust IR emission (in green). Compared to the SED of the full galaxy, the stellar component is modelled with a single Kurucz stellar atmosphere model for AB Aurigae. Additionally, the dust around AB Aurigae is much hotter than the dust in the diffuse ISM probed by the galactic SED. For this reason, its contribution to the NIR bump in the SED can be modelled with a Planck curve at 1600 K. The data also show the same spectral feature of the silicate peak at 10 μm, as well as a plateau at longer wavelengths, consistent with colder dust farther from the central star. The presence of a bump in the NIR, or in general of an excess of IR radiation compared to the stellar component alone, is the signature of dust being present around a young stellar object (YSO) and thus of a protoplanetary disk. Moreover, the evolution of the SED shows the different stages, or classes, of YSO, as dust settles in the midplane and accretes onto the central star or disperses over time via photoevaporation (Ercolano & Pascucci (2017), van der Marel (2023)).
By combining theoretical SED models with IR observations, we can infer properties of the observed dust. For instance, Draine et al. (2014) applied SED-fitting to study the spatial distribution of dust properties in the Andromeda Galaxy. By using *Spitzer* and *Herschel* IR data, they were able to reconstruct spatially-resolved maps of dust surface density, dust-to-gas ratio, and other dust-related properties. For instance, Fig. 1.5 shows the dust mass surface density map and radial profile. While the azimuthal direction seems to be roughly homogeneous, the dust surface density shows a distinct peaked profile in the radial direction, with a clear peak at $R \simeq 11.2$ kpc, which corresponds to a peak in H_I 21 cm emission, and thus of star formation.

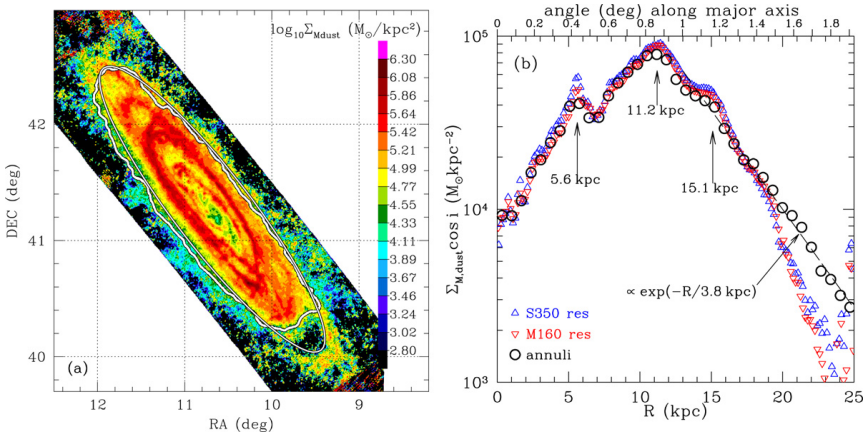

Figure 1.5: (a) Map of the dust mass surface density of M31 observed with *Spitzer* and *Herschel*. (b) Radial profile of the dust mass surface density. Image from Draine et al. (2014).

In terms of dust-to-gas ratio, Draine et al. (2014) find a decreasing trend from a central value of $M_{\rm d}/M_{\rm H} \approx 0.028$ to 0.005 at $R \approx 20$ kpc. This decreasing trend is consistent with the decreasing metallicity observed in similar spiral galaxies.

1.1.3 Presolar grains and dust composition

An additional and invaluable source of data regarding astophysical dust is provided by the dust which resides inside the Solar System. In particular, meteorites are a constant source of dust which is delivered on the surface of Earth. The most important class of meteorites, when speaking about astrophysical dust, are chondrite meteorites. These are the class of meteorites with similar chemical composition of the Sun, which come from undifferentiated asteroids. These are asteroids which haven't melted, so the chemical composition is still spread across the asteroid. Differentiated asteroids, on the other hand, present an heavier core rich in Fe and a rocky mantle. Chondritic meteorites are more primitive, and can tell the story of the evolution of material from the solar nebula to the present day.
Chondritic meteorites are composed of mainly three components:

- *Matrix*: represents the bulk of the meteoritic material. It is a collection of fine grained minerals and amorphus rocks, with inclusions of presolar grains and organic carbonaceous material.

- *Refractory inclusions*: irregular inclusions composed of high-temperature minerals. This is the temperature at which, from the gas phase, the molecules that form these inclusions have condensed in the solid phase. Being these temperatures very high ($T \gtrsim 1300$ K), these inclusions have usually formed during the initial stages of star formation, when temperatures in the disk were still high. As such, they are the oldest samples we have access to. Additionally, their composition is compatible with the expected composition of solids formed by cooling of solar composition gas. The most common type of refractory materials are calcium-aluminum-inclusions (CAIs).

- *Chondrules*: 0.1-1 mm spheres of silicates, probably formed as molten droplets and cooled in space.

Fig. 1.6 shows a chondritic meteorite with chondrules, matrix, and CAI clearly visible. Additionally, presolar grains found in the matrix are shown as enlargements (Nittler & Ciesla, 2016).
Apart from meteorites, other sources of presolar dust grains are interplanetary dust particles (IDPs) and spacecraft-returned samples. IDPs are collected by aircraft in the stratosphere. They are ~ 50 μm particles with a probable cometary origin. Spacecraft-returned

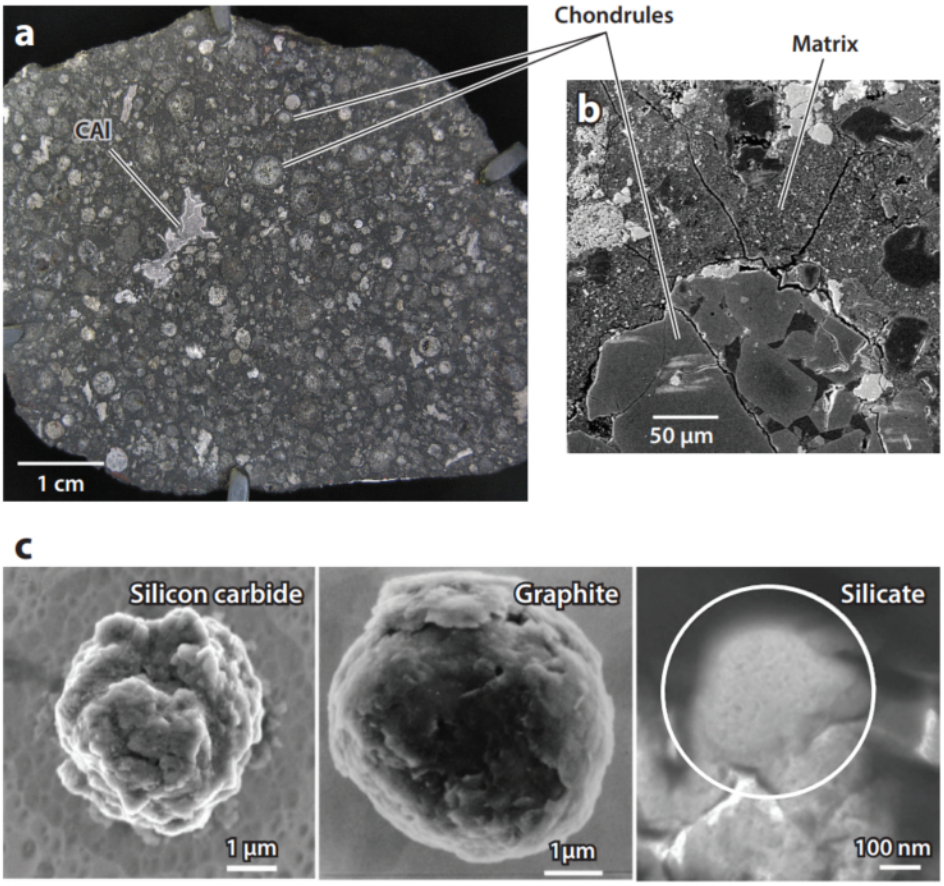

Figure 1.6: (a) Photograph of a carbonaceous chondrite showing a CAI, as well as chondrules. (b) Scanning electron microscope image of a carbonaceous chondrite, showing both the fine-grained matrix and chondrule inclusions. (c) Electron microscope images of presolar grains of different composition. Image from Nittler & Ciesla (2016).

samples are, for instance, objects collected by the Stardust mission on the Wild-2 comet, or the Hayabusa mission on the Itokawa asteroid.
Hoppe & Zinner (2000) classified the types of presolar grains found in chondritic meteorites. It is to be noted that the presolar grains found in meteorites have survived the high temperatures needed to form meteorites in the first place. Thus, their relative abundances are not representative of the grains that populate the diffuse ISM or molecular clouds. However, from the classification of Hoppe & Zinner (2000) we can infer two classes of interstellar dust: silicates and carbonaceous.
These presolar grains can be studied to identify the possible solar sources from which they formed. In particular, different sources yield different isotopic compositions. Fig. 1.7 shows the different isotopologues of oxygen that are formed by different solar sources, such as AGB stars, Novae, and Supernovae explosions (SNe) (Nittler & Ciesla, 2016). At the center of the Figure, the red square shows the average isotopologue composition of the

Solar System bodies. These presolar grains show a markedly different composition from that of the present-day Solar System, confirming their presolar origin. Additionally, the most common sources of these presolar grains seem to be AGB stars. This is compatible with the fact that, as we will see in Section 1.2, AGB stars are among the most efficient solar sources of astrophysical dust.

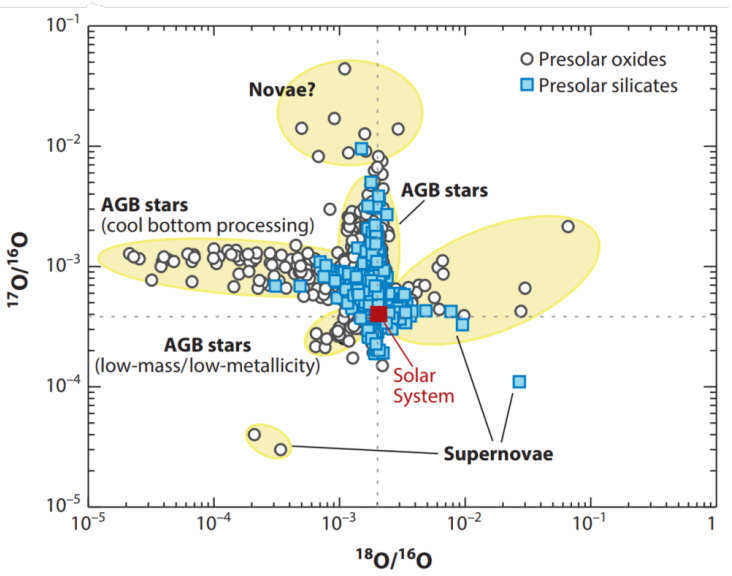

Figure 1.7: Oxygen isotopic ratios measured in presolar dust grains. The ratios span large ranges, deviating from Solar System values. The different clusters are marked by their different inferred stellar sources. Image from Nittler & Ciesla (2016).

1.2 The life cycle of dust

This section provides an introductory overview of the physical processes governing the life cycle of cosmic dust, from its formation in stellar environments to its evolution in the diffuse ISM, and finally through the stages of star and planet formation. Particular emphasis is placed on how the fundamental properties of dust grains evolve across these different environments.

Dust produced by stellar sources such as asymptotic giant branch stars and SNe is injected into the ISM with a GSD that differs markedly from the distribution inferred from observations of interstellar extinction and emission. In particular, the initial stellar GSD is generally biased toward larger grains and lacks the characteristic power-law behavior of the MRN distribution. This discrepancy implies that dust must undergo substantial processing in the diffuse ISM through mechanisms such as shattering, coagulation, accretion of gas-phase metals, and destruction by interstellar shocks. These processes collectively reshape the GSD and determine the optical and thermal properties of interstellar dust.

The section also briefly introduces the topic of dust evolution in protoplanetary disks,

where the physical conditions differ dramatically from those in the diffuse ISM. In disks, dust grains experience rapid growth through coagulation, radial transport driven by aerodynamic drag, and fragmentation in high-velocity collisions. These processes control the maximum grain size and spatial distribution of solids, and ultimately set the initial conditions for planetesimal formation. Together, these evolutionary stages illustrate how dust acts as a dynamic and evolving component throughout the star and planet formation cycle.

1.2.1 Dust formation

Dust Nucleation

Before discussing the astrophysical sources of dust, I briefly review the models used to describe the microscopic buildup of dust grains from individual monomers. This process, known as nucleation, is fundamental to understanding the conditions under which atoms and molecules can condense to form the initial dust seeds that subsequently grow to sizes ranging from ∼nm to μm. The simplest and most widely used framework is Classical Nucleation Theory (CNT), which treats the problem from a thermodynamic perspective. The free-energy difference required to condense a spherical grain of radius r

$$\Delta G = 4\pi r^2 \sigma - N k_{\mathrm{B}} T \ln \frac{P}{P_{\mathrm{sat}}}, \tag{1.13}$$

where σ is the surface tension, N is the number of monomers in the grain, k_{B} is the Boltzmann constant, P the pressure of the gaseous environment and P_{sat} the corresponding saturation pressure.
The first term, proportional to $\propto r^2$, represents the positive surface-energy contribution that suppresses condensation. The second term is a bulk contribution proportional to the grain volume through $N \propto r^3$. The combination of these competing terms determines the behavior of the free energy, as illustrated in Fig. 1.8. For radii $r < r_c$, the free energy increases, creating a barrier that must be overcome. For $r > r_c$, the negative slope means nucleation becomes energetically favored relative to evaporation. This occurs only when the bulk term in Eq. 1.13 is positive, that is, when $P > P_{\mathrm{sat}}$. This condition, known as supersaturation, is required for nucleation to proceed.
The application of CNT to dust-grain formation has been questioned because of several underlying assumptions. First, it is a thermodynamic theory that requires chemical equilibrium, which is generally absent in astrophysical environments because of their low monomer densities. Moreover, CNT cannot account for shocks or other non-equilibrium conditions. To overcome these limitations, alternative nucleation frameworks have been developed, most notably Kinetic Nucleation Theory (KNT). KNT does not assume thermodynamic equilibrium; instead, it models condensation and evaporation using a kinetic approach based on a master equation of the form

$$\frac{df_n}{dt} = \beta_{n-1} f_{n-1} - \beta_n f_n + \alpha_{n+1} f_{n+1} - \alpha_n f_n, \tag{1.14}$$

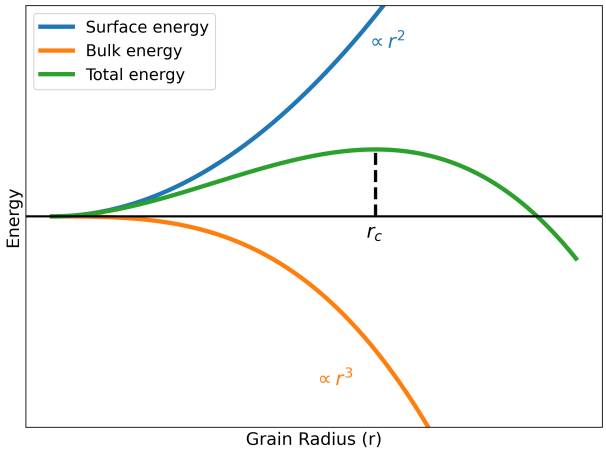

Figure 1.8: Schematic illustration of the interplay between surface and bulk energy in determining the free energy difference required for nucleation.

where f_n is the number density of clusters of size n, while α and β are the evaporation and condensation rates, respectively.
These are "top-down" approaches, in which microscopic nucleation behavior is derived from macroscopic thermodynamic input. The opposite, "bottom-up," approach has also been applied to astrophysical dust: direct quantum-chemical simulations of the dynamics of individual atoms. Although computationally far more expensive than CNT or KNT, these methods are better suited to the earliest stages of nucleation, when the number of monomers is $\lesssim 10$. Such models typically predict lower formation rates (Mauney & Lazzati, 2018).

AGB stars

The dense atmospheres of asymptotic giant branch (AGB) stars have traditionally been considered the environments in which most astrophysical dust is produced (Valiante et al., 2009; Di Criscienzo et al., 2013; Schneider & Maiolino, 2024). Low- and intermediate-mass stars ($0.8M_\odot < M_{\rm ZAMS} < 8M_\odot$, where $M_{\rm ZAMS}$ is the zero-age-main-sequence star mass) end their lives in this phase of strong mass loss. During this stage, pulsations of the stellar interior generate shock waves that propagate through the outer layers, ejecting material. The dense and relatively cool envelopes produced in this way provide favorable conditions for dust condensation.
The types of grains that form in these environments depend strongly on the C/O ratio. Because CO is a very stable molecule, if C/O < 1, most carbon atoms are locked into CO and are therefore unavailable for forming carbonaceous grains. Oxygen remains available to form SiO and, if the metallicity is sufficiently high, $MgSiO_3$ and Mg_2SiO_4, which subsequently condense into silicate grains. Conversely, if C/O > 1, essentially all oxygen is bound in CO, suppressing silicate production and favoring carbonaceous dust.

Fig. 1.9 shows the dust masses predicted for AGB stars of various masses and metallicities for both carbonaceous and silicate grains, based on Zhukovska et al. (2008). The resulting dust properties depend on both stellar mass and metallicity. Stars more massive than $4M_{\odot}$ undergo hot bottom burning (HBB), a process in which the base of the convective envelope overlaps with the hot hydrogen-burning shell, enabling nuclear processing within the envelope itself. HBB can drive CN-cycle reactions that convert carbon into nitrogen, reducing the C/O ratio. As discussed above, decreasing C/O suppresses carbonaceous dust formation and promotes silicate production, as seen in the left panel of Fig. 1.9. Lower-mass AGB stars do not experience HBB, so they typically maintain C/O > 1, leading to the opposite trend.

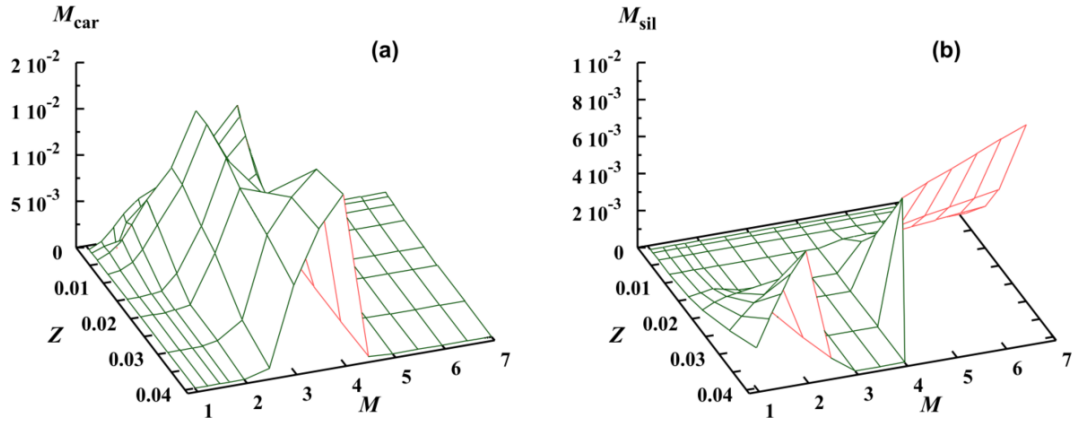

Figure 1.9: Dust masses obtained from simulated AGB stars of different masses and metallicities for (a) carbonaceous grains and (b) silicate grains. All masses are in units of $M_{\odot}$. Adapted from Zhukovska et al. (2008).

Metallicity also plays a significant role. Carbonaceous dust shows little dependence on metallicity, while silicate grains are produced only when the metallicity is high enough to supply adequate Mg for forming $MgSiO_3$ and Mg_2SiO_4. Using SED fitting of AGB spectra, Groenewegen (1997) inferred relatively large grain sizes, on the order of 0.1 μm.
Dust not only is formed in the outer envelopes of AGB stars but also plays a crucial dynamical role in the mass-loss process (Höfner & Olofsson, 2018). AGB stars exhibit large mass-loss rates, approximately 10^{-5} $M_{\odot}yr^{-1}$, driven by winds with velocities of about 10 kms^{-1}. These winds are powered by radiation pressure from the central star. Material is accelerated outward when the radiative acceleration exceeds the gravitational pull:

$$\frac{a_{\mathrm{rad}}}{a_{\mathrm{grav}}} = \frac{\langle \kappa_{\lambda} \rangle L_{*}}{4\pi c G M_{*}}, \tag{1.15}$$

where $\langle \kappa_{\lambda} \rangle$ is the average opacity, regulating how efficiently photons are absorbed, L_{*} and M_{*} are the luminosity and mass of the central star, c is the speed of light and G the gravitational constant.

Because AGB stars have low effective temperatures of around 3000 K, their blackbody emission peaks near 1 μm in the near-infrared (NIR). At these wavelengths, atomic transitions are too inefficient to drive strong winds. Molecular roto-vibrational transitions can absorb NIR photons but appear to be too weak to account for the observed outflow strengths. Dust grains, however, can efficiently absorb NIR radiation, accelerate outward, and entrain the surrounding gas.

Type II SNe

Type II SNe are believed to be among the earliest sources of dust, owing to the short lifetimes of their massive progenitors (Hirashita, 2013). The left panel of Fig. 1.10 shows the SEDs of SN2004et observed with the IRAC camera aboard Spitzer over time Kotak et al. (2009). These SEDs reveal signatures of dust formation and evolution. After about 300 days post-explosion, the spectra exhibit narrow-band features of SiO molecules and broad-band emission from silicate dust. Over time, the SiO features weaken and eventually disappear by day 795, possibly indicating that SiO molecules have been incorporated into newly formed silicate grains (Cherchneff, 2014).

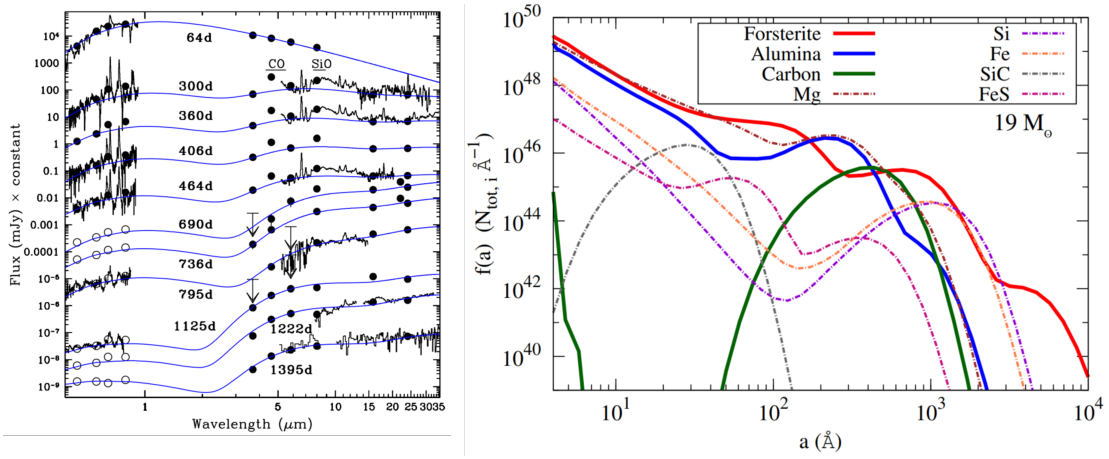

Figure 1.10: (a) SEDs of SN2004et from *Spitzer* observations for different days after the SNe. Image from (Kotak et al., 2009). (b) GSDs of different dust components after a $19 M_{\odot}$ SNe. The formation and evolution of dust grains are modelled using CNT. Image from Sarangi & Cherchneff (2015).

Models of dust nucleation in SNe ejecta support the formation of dust in this environment, particularly during the adiabatic expansion phase. The right panel of Fig. 1.10 shows the resulting grain-size distributions (GSDs) for different dust species obtained by applying CNT to a 19 $M_{\odot}$ SNe (Sarangi & Cherchneff, 2015). The GSD decreases with increasing grain size and reaches maximum radii of approximately $a \approx 0.1 - 1$ μm.
However, grains produced in these early phases must still survive transport into the ISM. In particular, dust needs to withstand the reverse shock generated as the outward-moving blast wave interacts with the ambient medium. Two key destructive processes govern dust

evolution in this environment: sputtering and shattering. Sputtering refers to the erosion of dust grains by energetic gas particles that strike and eject surface atoms. Because it is a surface process, sputtering is more efficient for small grains than for large ones. Shattering, in contrast, is the fragmentation of grains through grain–grain collisions.

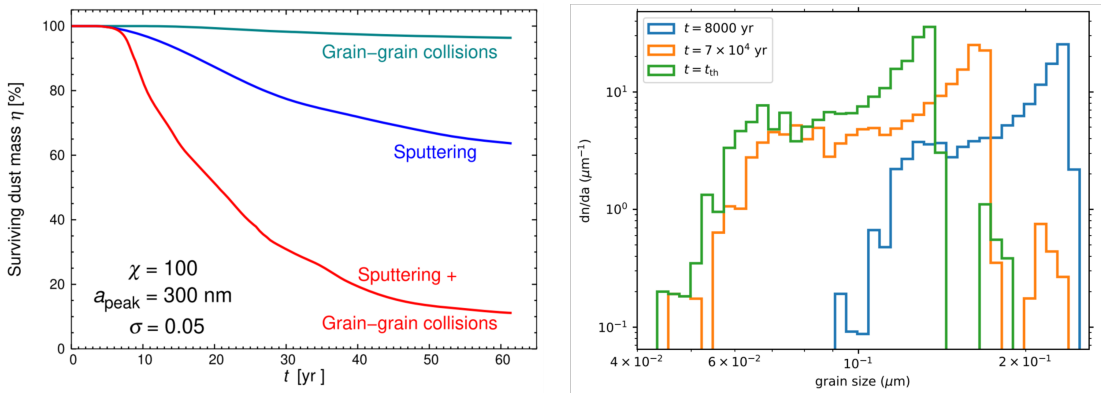

Figure 1.11: (*a*) Fraction of dust mass surviving after passage through a shock, shown as a function of time. Curves include destruction by grain–grain collisions (green), sputtering (blue), and both processes combined (red). Because the two mechanisms act synergistically, the combined destruction exceeds the sum of the individual effects Kirchschlager et al. (2019). (*b*) Evolution of the GSD of silicate grains after crossing the reverse shock at $t = 8000\ yr$. Continued sputtering by the ISM erodes large grains into smaller ones until sputtering becomes inefficient at $t = t_{\rm th}$. Image from Slavin et al. (2020).

The left panel of Fig. 1.11 shows the effects of sputtering, shattering, and their combination on the surviving dust mass in a shock-processed region (Kirchschlager et al., 2019). A key result is that sputtering and shattering act synergistically: together, they destroy more dust than the sum of their individual contributions. Shattering breaks large grains into smaller fragments, making them more vulnerable to sputtering. When both processes are included, only about $\sim 10\%$ of the initial dust mass survives.
After passing through the reverse shock, the remaining grains have been significantly accelerated and must decelerate while interacting with the ISM. During this deceleration phase, sputtering continues to erode the grains. The right panel of Fig. 1.11 presents results from 2D hydrodynamical simulations that include sputtering (Slavin et al., 2020). At $t = 8000$ yr, the grains have just crossed the reverse shock; at this stage, the GSD is strongly biased toward large radii because small grains have been efficiently destroyed. Subsequent interaction with the ISM steadily shifts the GSD toward smaller radii, as shown at $t = 7{\cdot}10^4$ yr and later at $t = t_{\rm th}$, when the grains have slowed sufficiently that sputtering becomes inefficient.
Despite this further processing, the final GSD remains skewed toward large grains, opposite to the MRN-like GSD observed in the diffuse ISM. To understand how the observed ISM GSD arises from the dust contributed by AGB stars and SNe, we now investigate how dust is processed in the diffuse ISM.

1.2.2 Dust processes in the diffuse ISM

Once dust is injected into the ISM by stellar sources such as AGB stars and SNe, its life cycle has only just begun. As discussed at the end of the previous section, the grain-size distributions (GSDs) produced by these stellar sources are significantly more top-heavy than the observed MRN distribution. Additional processing is therefore required to reconcile the initial grain population with observations of the diffuse ISM. This section highlights several key dust processes that operate in the diffuse ISM.

Dust destruction by SNe shocks

Dust grains in the diffuse ISM experience significant mass loss when they encounter SNe shocks. Once the ISM is shocked, small grains ($a \lesssim 0.01$ μm) become entrained in the hot, post-shock gas and remain trapped there, while larger grains tend to decouple from the gas and stay in the cooler post-shock region. As a result, small grains are efficiently destroyed by sputtering, whereas the population of larger grains is less affected.
Sputtering is the erosion of dust grains through interactions with energetic atoms and ions. It occurs in two forms: thermal and non-thermal. Thermal sputtering takes place when dust grains are immersed in hot gas and are eroded by high-temperature particle impacts. Non-thermal sputtering occurs when dust and gas have a significant relative velocity, so that collisions are ballistic rather than thermal. Both types of sputtering can be described by appropriate erosion rates: the thermal rate, $Y_{\rm th}(T)$, which depends on the gas temperature, and the non-thermal rate, $Y_{\rm non\text{-}th}(v_{\rm rel})$, which depends on the relative velocity between gas and dust.
The resulting mass-loss rate due to sputtering can be written, following Hu et al. (2019), as

$$\frac{dm_{\rm d}}{dt} = \frac{3n_{\rm H}m_{\rm d}}{a}Y_{\rm tot}, \tag{1.16}$$

where $n_{\rm H}$ is the hydrogen number density, $m_{\rm d}$ the dust mass, and $Y_{\rm tot} = Y_{\rm th}(T)+Y_{\rm non\text{-}th}(v_{\rm rel})$ is the total erosion rate. Because sputtering is a surface process, it is more effective for small grains than for large ones.
This mass-loss rate can be recast as a sputtering timescale:

$$\tau_{\rm sp} = \frac{a}{3n_{\rm H}Y_{\rm tot}} \simeq 0.33{\rm Myr}\left(\frac{a}{\mu{\rm m}}\right)\left(\frac{n_{\rm H}}{cm^{-3}}\right)^{-1}\left(\frac{Y_{\rm tot}}{10^{-6}\ \mu myr^{-1}cm^{3}}\right)^{-1}. \tag{1.17}$$

This timescale is quite short compared to the dynamical timescale of a SNe remnant, indicating that sputtering in SN shocks is highly efficient at destroying small grains embedded in hot gas.
We can also estimate a characteristic timescale for dust destruction across the entire Milky Way:

$$\tau_{\rm SN} = \frac{M_{\rm g}}{\epsilon m_{\rm swept}\gamma_{\rm SN}}, \tag{1.18}$$

where $M_{\rm g}$ is the total ISM gas mass, $\epsilon \sim 0.1$ is the dust-destruction efficiency, $m_{\rm swept}$ is the mass of gas swept up by a single SNe, and $\gamma_{\rm SN}$ is the SNe rate. In the Milky Way,

$\tau_{\rm SN} \sim 4 \times 10^8$ yr, which is much shorter than the timescale for dust production by stellar sources. The latter is roughly comparable to the gas-consumption timescale, $\tau_{\rm SF} \sim M_{\rm g}/\Psi$, where Ψ is the star-formation rate. For the Milky Way, $\tau_{\rm SF} \sim 3 \times 10^9$ yr.
Since the destruction timescale is much shorter than the stellar dust-production timescale, additional dust-formation and growth processes must be operating directly within the ISM.

Gas accretion

Dust grains can grow by directly accreting gas-phase metals onto their surfaces. For accretion to be effective, the environment must be sufficiently cold; therefore, gas accretion operates primarily in the cold neutral medium and in molecular clouds. Here, we derive an expression for the accretion timescale, following Asano et al. (2013), in order to compare it with both the destruction timescale from SN shocks and the production timescale from stellar sources.
We begin by expressing the dust mass growth rate as

$$\left(\frac{dM_{\rm d}}{dt}\right)_{\rm acc} = \eta N_{\rm d} \pi \langle a^2 \rangle \alpha \rho_{Z,{\rm g}} \langle v_{\rm th} \rangle, \tag{1.19}$$

where η is the mass fraction of clouds, $N_{\rm d}$ the number of dust grains, $\langle a^2 \rangle$ is the second moment of the grain size, α is the mean sticking coefficient of metals on the dust surface, $\rho_{Z,{\rm g}}$ is the mass density of metals in the gas phase and $\langle v_{\rm th} \rangle$ is the mean thermal velocity of the gas.
Assuming spherical grains, we can express the grain number in terms of the total dust mass and the mass of individual dust grains:

$$N_{\rm d} = \frac{M_{\rm d}}{m_{\rm d}} = \frac{3M_{\rm d}}{4\pi \langle a^3 \rangle \rho_{\rm gr}}, \tag{1.20}$$

where $\rho_{\rm gr}$ is the internal mass density of the grains. The gas-phase metal density can be written in terms of the metallicity:

$$\rho_{Z,{\rm g}} = \rho_{\rm ISM}^{\rm eff} Z (1 - \delta), \tag{1.21}$$

where $\rho_{\rm ISM}^{\rm eff}$ is the mass density of material in interstellar clouds, Z the metallicity, and δ the fraction of metals already locked in dust.
Combining these expressions, the accretion timescale becomes

$$\tau_{\rm acc} = \frac{M_{\rm d}}{dM_{\rm d}/dt} = \frac{4\langle a^3 \rangle \rho_{\rm gr}}{3\eta \langle a^2 \rangle \alpha \rho_{\rm ISM}^{\rm eff} (1-\delta) \langle v_{\rm th} \rangle} \cdot \frac{1}{Z} \tag{1.22}$$

$$\approx 2.0 \times 10^7 \; yr \left(\frac{\overline{a}}{0.1 \mu m}\right) \left(\frac{n_{\rm H}}{100 cm^{-3}}\right)^{-1} \left(\frac{T}{50 \; K}\right)^{-\frac{1}{2}} \left(\frac{Z}{0.02}\right)^{-1} = \tau_{\rm acc,0} Z^{-1}, \tag{1.23}$$

where $\tau_{\rm acc,0} = 4 \times 10^5 \; yr$.
This accretion timescale is much shorter than both the dust destruction timescale from SN

shocks and the dust production timescale from stellar sources, indicating that accretion is a strong candidate to explain the net growth of dust grains in the ISM. Because the timescale depends inversely on metallicity, there exists a critical metallicity, $Z_{\rm crit}$, at which accretion becomes the dominant growth mechanism. Asano et al. (2013) found $Z_{\rm crit} \approx 0.1\ Z_{\odot}$ for a Milky Way–like galaxy. Such a low value implies that accretion dominates under most ISM conditions. Furthermore, because $\tau_{\rm acc}$ decreases as the metallicity increases, accretion is a runaway process: once it becomes dominant, it rapidly enhances the dust mass fraction.

Grain-grain collisions

Grain–grain collisions constitute another important mechanism that modifies the GSD. These collisions do not change the total dust mass, but they can redistribute it: small grains can stick together to form larger aggregates (coagulation), or large grains can fragment into smaller pieces (shattering).
The statistical evolution of the GSD under collisions is described by the Smoluchowski equation (Birnstiel, 2024). One convenient form is

$$\frac{dn(m)}{dt} = \int_0^{\infty}\int_0^{m_1} K(m|m_1,m_2)R(m_1,m_2)n(m_1)n(m_2)dm_1dm_2 \tag{1.24}$$

$$- n(m)\int_0^{\infty} R(m,m_1)n(m_1)dm_1, \tag{1.25}$$

where $n(m)$ is the number density of grains with mass m, $R(m_1,m_2)$ is the collisional rate of two particles of mass m_1 and m_2, and the $K(m|m_1,m_2)$ is the coagulation kernel, describing the conditional probability of ending up with a grain of size m from the interaction of two grains of sizes m_1 and m_2. The first term encodes the growth of $n(m)$, either by fragmenting larger grains or by cogulating smaller ones. The second term, on the other hand, encodes the mass loss of $n(m)$ by interactions of grains of mass m with all other mass bins.
The collision rates $R(m_1,m_2)$ can be written in their simplest form as

$$R(m_1,m_2) = \sigma_{12}\langle\Delta v_{12}\rangle \quad \sigma_{12} = \pi(a_1+a_2)^2, \tag{1.26}$$

where σ_{12} is the geometrical cross-section and $\langle\Delta v_{12}\rangle$ is the average relative velocity between the two dust mass bins. As for the conditional probabilities $K(m|m_1,m_2)$, they encode all possible collisional outcomes and encode the real complexity of the problem. For the simple case in which all collisions lead to perfect sticking, they can simply be modeled by Dirac delta functions

$$K(m|m_1,m_2) = \delta(m-m_1-m_2). \tag{1.27}$$

In reality, collisional outcomes are more complex: sticking, bouncing, erosion (cratering), and catastrophic fragmentation can all occur depending on particle properties and impact conditions. Fig. 1.12 shows experimental collisional outcomes for dust aggregates of different sizes and for different relative velocities. Small particles with small relative velocities usually end up sticking and forming larger grains. On the opposite end of the parameter

space, large grains with large relative velocities fragment and produce smaller grains. As can be seen in Fig. 1.12, sticking and fragmentation are not the only possible outcomes. Dust grains of the same size can bounce back upon collision, while grains of different sizes can experience erosion or cratering.

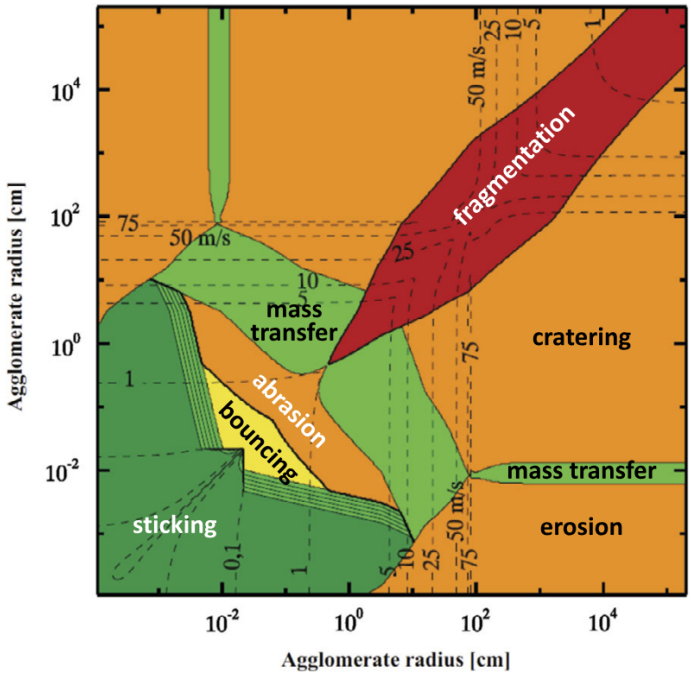

Figure 1.12: Collisional outcomes determined experimentally for pairs of dust aggregates, depending on their sizes and relative velocities. Credits: Stefan Kothe and Jürgen Blum.

Even if Fig. 1.12 is only illustrative of the possible outcomes of dust collisions, it represents quantitatively also what is believed to happen in the diffuse ISM. Large grains are accelerated by MHD turbulence (Yan et al., 2004) and end up having large velocity dispersions, and thus large relative velocities. For grains $a \gtrsim 0.1\mu m$, the velocity dispersion can be as high as $\sigma_v \sim 10 kms^{-1}$, causing fragmentation upon collision. Smaller grains, instead, maintain smaller relative velocities and favour coagulation instead. Fig. 1.13 shows the evolution of the GSD in a Milky Way galaxy for pure coagulation and pure shattering (Hirashita et al., 2021). The initial condition is set as the MRN distribution, which is then evolved under pure coagulation (left) or pure shattering (center). Coagulation shifts the GSD towards larger grain sizes, depleting the small-grain population in favour of large grains. Shattering does the opposite, reducing the number density of large grains while populating the smaller mass bins. The timescale for these processes found by Hirashita et al. (2021) is of the order of $\tau_{\rm gg} \sim 10^8\ yr$.
When both processes are combined, the evolution maintains the GSD close to the initial MRN distirbution, as shown in the right panel of Fig. 1.13. This suggests that the observed distribution in the diffuse ISM is the equilibrium distribution for dust under coagulation

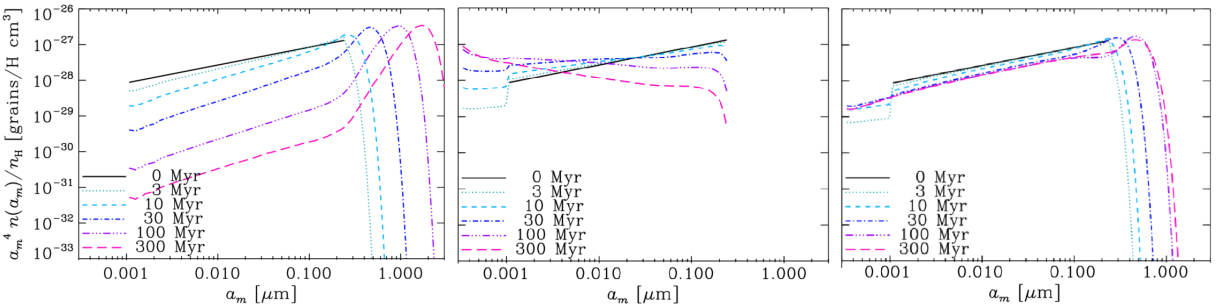

Figure 1.13: Evolution of the GSD due to pure coagulation (left), pure shattering (center), or both (right), starting from an MRN-like GSD. Figure adapted from Hirashita et al. (2021).

and fragmentation.

1.2.3 Back to the stars: the role of dust in star and planet formation

In this section, we close the life cycle of dust with its return to the stars. After being delivered to the ISM and being reprocessed, dust plays a significant role during star and planet formation. The fate of dust will be that of being accreted by the star, so that the metals of which it is composed can be recycled when the star ends its life and returns its materials into the ISM. However, dust in protoplanetary disks may first assemble, grow, and become pebbles, planetesimals, and later, planets.

Dust from Molecular Clouds to Protoplanetary Disks

Molecular clouds are approximately in hydrostatic equilibrium, where the inward pull of gravity is balanced by outward thermal pressure. For a region within the cloud to collapse and form a star, this balance must be upset: the gas must cool efficiently enough that gravitational attraction overwhelms thermal pressure. A useful criterion for collapse is the Jeans mass: the minimum mass that becomes gravitationally unstable at a given density and temperature. In its simplest form, the Jeans mass scales as

$$M_J \propto T^{\frac{3}{2}} n^{-\frac{1}{2}}, \tag{1.28}$$

where T is the gas temperature and n is the gas number density. Lower temperatures lead to a smaller Jeans mass and therefore favor the formation of low-mass, gravitationally bound fragments.

A molecular cloud can shed thermal energy via several radiative processes. A major contributor is line emission from atoms and molecules present in the gas. However, at high densities, the optical depths of these cooling lines increase rapidly. As a result, line photons

are frequently absorbed and re-emitted multiple times, so line cooling becomes inefficient deep inside dense cloud regions.
At this stage, dust becomes the dominant coolant. Dust grains can absorb thermal energy from the gas through collisions. Each collision transfers a small amount of kinetic energy from gas particles to the solid grain, exciting vibrational and rotational modes of the dust material. Dust grains then radiate this energy away very effectively in the IR. Crucially, the cloud is largely transparent to IR radiation, so these photons escape, carrying energy out of the dense, shielded interior.
Dust grains also influence cloud cooling indirectly by catalyzing the formation of molecular hydrogen. In cold, shielded environments, the gas-phase formation of H_2 is extremely inefficient, because gas-phase reactions require catalysts (e.g., free electrons in primordial gas) that are scarce in molecular clouds. Instead, dust grains act as catalytic surfaces:

- Physisorption: Hydrogen atoms from the gas stick to dust-grain surfaces through weak Van der Waals forces.

- Surface diffusion: Adsorbed H atoms migrate across the surface via thermal hopping or quantum tunneling, seeking binding sites.

- Reaction: When two H atoms meet, they react to form H_2. The reaction energy (4.48 eV) is partially deposited into the grain lattice.

- Desorption: The newly formed H_2 molecule typically desorbs from the surface with a modest kinetic energy.

The formation of H_2 aids cooling in two ways. First, it removes energetic H atoms from the gas, replacing them with cooler H_2 molecules that contribute less to thermal pressure. Secondly, H_2 itself is a coolant via rovibrational line emission.
Fig. 1.14 illustrates the Jeans mass predicted when only gas-phase line cooling is included versus when dust cooling is also accounted for (Schneider et al., 2006). Even at metallicities well below solar, dust dominates the cooling budget once the density is high enough for line cooling to saturate. As a result, fragments as small as $\sim 1\ M_\odot$ become possible. In contrast, if dust cooling were absent and only gas cooling operated, the minimum fragment mass would remain much larger, typically $\sim 10^2 M_\odot$.
It remains an open question whether dust undergoes significant evolution during the transition from the molecular cloud envelope, where dust is inherited directly from the diffuse ISM, to the protoplanetary disk. This stage of dust evolution is poorly constrained because protostellar cores are particularly challenging to observe and characterize.
Extinction-based observations, which work well on large spatial scales, become ineffective in the dense interiors of molecular clouds due to high optical depth. Conversely, thermal dust emission, used to probe protoplanetary disks, becomes too faint at large distances

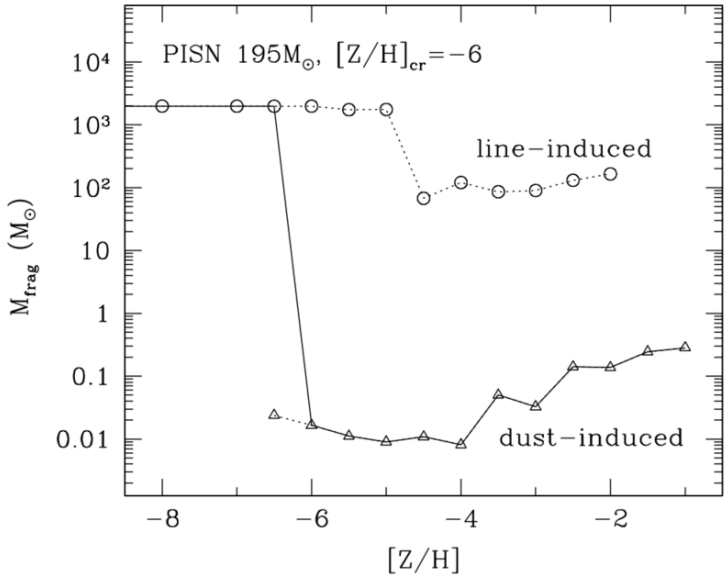

Figure 1.14: Jeans mass corresponding to line and dust-induced cooling mechanisms (circles and triangles, respectively). The solid line indicates the transition in fragment masses when dust-induced cooling becomes prominent, for metallicities $Z > 10^{-6}\ Z_{\odot}$. Image from Schneider et al. (2006).

from the central protostar. These observational limitations leave a gap in our ability to trace dust evolution continuously from ISM conditions down to disk scales.
Recent observations by Cacciapuoti et al. (2025) have helped narrow this gap by constraining the maximum dust-grain size through the spectral index, α, of the protostellar-core SED. The spectral index is defined observationally as

$$\alpha = \frac{\log F_{\nu_2} - \log F_{\nu_1}}{\log \nu_2 - \log \nu_1}, \tag{1.29}$$

where ν_1 and ν_2 are the two frequencies used to measure the slope of the SED.
Under the optically thin approximation (Eq. 1.12), the IR flux depends on the product of dust opacity and the Planck function. In the Rayleigh–Jeans limit, the Planck function scales as $B_\nu(T) \propto \nu^2$. The dust opacity typically follows a power-law behavior, $\kappa_\nu \propto \nu^\beta$, where the exponent β is large for small grains, which have steep emissivity spectra and radiate relatively more at short wavelengths, and decreases for larger grains that exhibit a flatter opacity curve.
Combining these trends, we obtain

$$F_\nu \propto \nu^\alpha = \nu^{2+\beta}. \tag{1.30}$$

Hence, by measuring α, we can infer β and thus estimate the characteristic grain size contributing to the emission. In the diffuse ISM, typical values are $\alpha \sim 3.5$.
Fig. 1.15 summarizes the results of Cacciapuoti et al. (2025). Protostellar envelopes exhibit

spectral indices intermediate between those of the diffuse ISM and protoplanetary disks, indicating that dust growth or structural modification may already be underway before disk formation.

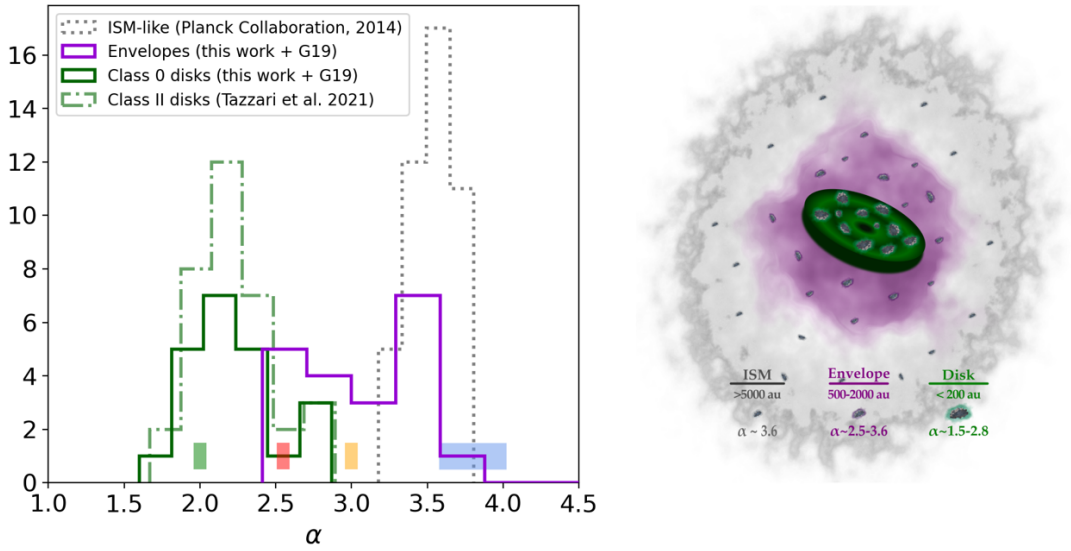

Figure 1.15: (*a*) Spectral index distributions measured at different spatial scales, from ISM values obtained by the Planck Collaboration to those observed in evolved protoplanetary disks. Protostellar envelopes show intermediate values, suggesting ongoing dust evolution during the earliest phases of star formation. (*b*) Schematic representation of dust evolution from the ISM to the protoplanetary disk. Image from Cacciapuoti et al. (2025).

Dust evolution in Protoplanetary Disks

Dust plays a central role in virtually every aspect of protoplanetary disk physics. Observationally, its IR emission is an invaluable probe for studying both dust properties via SEDs (Meeus, G. et al., 2001) as well as disk structure and morphology with spatially resolved ALMA observations (Andrews et al., 2018). Physically, dust regulates a wide range of dynamical and chemical processes and follows evolutionary pathways that either lead to its accretion onto the central star or its growth into larger bodies, the precursors of rocky planets and cores of gas giants.

Given the complexity of dust–gas interactions in disks, it is not possible to provide an exhaustive overview here. In this section, I focus on two key processes that control the evolution of dust in protoplanetary disks: radial drift and growth barriers. For a more comprehensive review, see Birnstiel (2024) and references therein.

Before discussing these processes, we introduce several concepts relevant to dust dynamics. The first is the stopping time, the characteristic timescale over which a dust grain adjusts its velocity relative to the gas. Its form depends on the aerodynamic regime, which is determined by whether the dust interacts with the gas as a continuous fluid or as a sequence

of individual molecular collisions. When the grain size is smaller than the gas mean free path ($a < \lambda_{\rm mfp}$), the dust is in the Epstein regime, and the stopping time becomes

$$t_{\rm s} = \frac{a\rho_{\rm gr}}{\langle v_{\rm th}\rangle \rho_{\rm g}}. \tag{1.31}$$

Intuitively, larger or denser grains have longer stopping times, while higher gas density and larger thermal velocity shorten $t_{\rm s}$.
A related dimensionless parameter is the Stokes number, defined as

$$\mathrm{St} = t_{\rm s}\Omega_K, \tag{1.32}$$

where Ω_K is the Keplerian orbital frequency. The Stokes number characterizes the coupling between dust and gas independently of the local physical scales.
Once dust has settled into the disk, gas and dust evolve under gravity, drag forces, and (for gas only) pressure forces. The resulting acceleration equations are

$$\frac{\partial \vec{v}_{\rm g}}{\partial t} = \frac{\vec{v}_{\rm d} - \vec{v}_{\rm g}}{t_{\rm s}}\epsilon - \frac{GM_*}{r^3}\vec{r} - \frac{\nabla P}{\rho_{\rm g}} \tag{1.33}$$

$$\frac{\partial \vec{v}_{\rm d}}{\partial t} = -\frac{\vec{v}_{\rm d} - \vec{v}_{\rm g}}{t_{\rm s}} - \frac{GM_*}{r^3}\vec{r}, \tag{1.34}$$

where $\epsilon = \rho_{\rm d}/\rho_{\rm g}$ is the dust-to-gas mass ratio. At equilibrium, dust moves on Keplerian orbits ($v_{{\rm d},\phi} = v_{\rm K}$), while gas orbits slightly more slowly. Gas experiences an outward pressure-gradient force, partially pushing outward against gravity and leading to a sub-Keplerian rotation speed,

$$v_{{\rm g},\phi} = (1-\eta)v_{\rm K}, \qquad \eta = -\frac{1}{2}\left(\frac{h_{\rm g}}{r}\right)^2 \frac{\partial \ln P}{\partial \ln r} \sim 10^{-3}\sqrt{r/\mathrm{AU}}, \tag{1.35}$$

where $h_{\rm g}$ is the gas scale height. Although the deviation from Keplerian rotation is small, it has major implications for dust dynamics.
Because dust orbits slightly faster than gas, it experiences a headwind that removes angular momentum and slows it down. To maintain centrifugal balance, the dust drifts inward to smaller radii, a process known as radial drift. This drift persists until the grain is accreted onto the star.
The efficiency of radial drift depends strongly on the Stokes number. Fig. 1.16 shows the radial velocities of gas and dust as a function of St. For small grains ($\mathrm{St} \ll 1$), strong coupling forces them to move nearly with the gas, leading to minimal drift. Large grains ($\mathrm{St} \gg 1$) are only weakly affected by the headwind and remain close to Keplerian orbits. Maximum drift occurs at St=1, where grains migrate inward at approximately $\eta v_{\rm K}$. The figure also illustrates how this behavior changes when including the backreaction of dust on gas. In the limit $\epsilon = 1$, dust can push the gas toward larger radii, producing opposite radial motions.
From this simple picture, it seems that all dust will eventually radially drift towards the

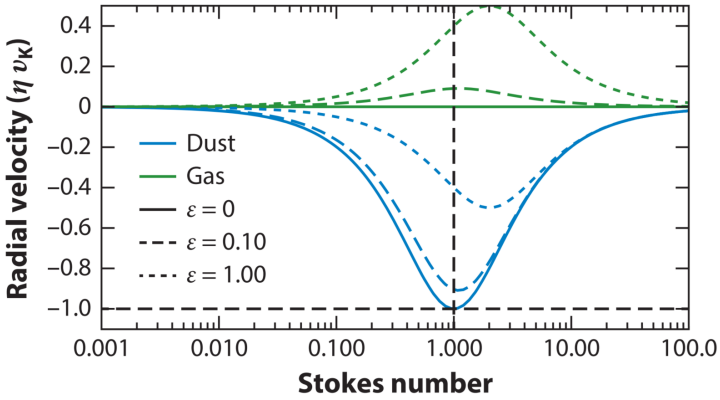

Figure 1.16: Dust (blue) and gas (green) radial velocities as a function of the Stokes number of the dust grains. The solid lines represent a case with no backreaction onto the gas, whereas the dashed lines show the effect of backreaction for dust-to-gas ratios of $\epsilon = 0.1$ and $\epsilon = 1.0$. Adapted from Birnstiel (2024), with data from Nakagawa et al. (1986).

star, thereby extinguishing all the solid material needed to form planets. However, this process can be steered towards a different outcome by the presence of a pressure bump in the disk. The consequence of the radial drift is that dust will move towards the maximum of the gas pressure profile. For a smooth disk, the maximum is reached at the center, so the dust will move only inward. If a pressure bump forms somewhere in the disk, the dust will tend to move toward it and accumulate at the pressure maximum. Different mechanisms have been proposed to generate pressure maxima: turbulence, vortices, infall of material from the envelope, or the presence of other planets already formed. The hypothesis that dust is trapped in pressure bumps is strengthened by spatially-resolved observations of dust IR emission by ALMA, which revealed substructures and rings in many of the observed protoplanets.

Dust in disks also undergoes coagulation and fragmentation. As grains grow, their Stokes number increases, making radial drift more efficient and preventing continuous growth. Radial drift thus acts as a growth barrier, limiting the maximum size of grains that can survive in a smooth disk.

Even if radial drift is suppressed by pressure traps, fragmentation imposes another growth barrier. As shown in Fig. 1.12, collisions between sufficiently large grains lead to fragmentation rather than growth. Fig. 1.17 summarizes how these processes shape the GSD.

Panel (a,i) shows the trajectories of dust grains in the radius-particle size space for a smooth disk. For each point in the grid, the green arrows show the amount of growth and radial drift dust experiences, and the black arrow shows the combined effect. We can see, for example, how larger grains drift faster with respect to smaller ones. The colormap shows the dust surface density per decade in size $\sigma(r, a)$. In red, the trajectories of specific grain parcels are shown, starting at the bottom of the GSD at different disk radii. As particles grow in size, they migrate inward, as expected by the radial drift theory discussed

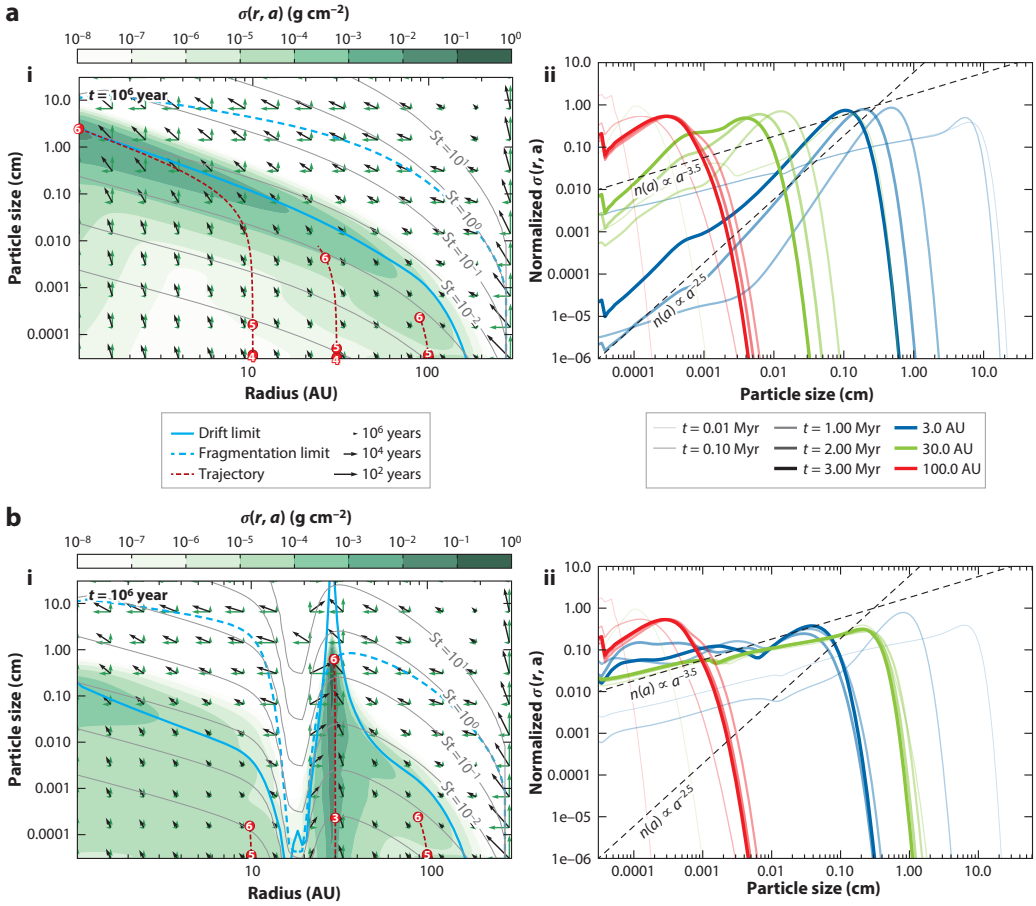

Figure 1.17: Evolution of dust particle size and radial position for a smooth disk (panels a) and in the presence of a pressure bump (panels b) from a Jupiter-mass planet at 20 AU. Panels (a,i) and (b,i) show the full GSD at all radii. The arrows represent the direction of radial motion and size growth. The red dashed lines and circles show the trajectories of some dust particles, where the numbers represent the time needed to reach that point (for instance, the number "5" indicates a time of 10^5 yr). The solid blue line indicates the radial drift barrier, whereas the dashed blue line indicates the fragmentation barrier. Panels (a,ii) and (b,ii) show slices of the GSD at 3 AU, 30 AU, and 100 AU. Figure from Birnstiel (2024).

above. At some point, the dust will reach the drift limit, the solid blue line, after which it will quickly migrate inward and be accreted onto the central star. Panel (*a*,*ii*) illustrates how the GSD evolves at different radii. At, for example, 3 AU, an initially MRN-like distribution $n(a) \propto a^{-3.5}$, becomes rapidly skewed toward larger sizes due to coagulation, but the maximum grain size is limited to $\sim$ 1 cm because larger grains drift inward too rapidly.

In the presence of a pressure bump, the situation changes markedly. Panel (*b*,*i*) shows that grains remain trapped and can grow to larger sizes than in the drift-limited case. Growth eventually reaches the fragmentation limit (dashed line), where collisions prevent further size increase. Panel (*b*,*ii*) shows that the GSD at 30 AU, where the pressure bump is located, maintains a near steady-state shape governed by the balance between coagulation and fragmentation. The total dust mass also remains roughly constant within the trap. Even in this case, however, grains do not grow beyond $\sim$ 1 cm.

Since standard disk evolution cannot easily overcome these growth barriers, additional processes are required to bridge the gap between dust grains and gravitationally bound planetesimals (Drążkowska et al., 2023). Proposed mechanisms include porous growth (Kataoka et al., 2013), and the streaming instability (Youdin & Goodman, 2005; Lesur et al., 2023).

2

Dust dynamics

In this chapter, we provide a concise overview of the theoretical framework used to describe the dynamics of dust in astrophysical environments. We begin from a kinetic description based on the Boltzmann equation and outline the sequence of modelling assumptions that lead to a continuum formulation in which dust is treated as a fluid. In particular, we discuss the physical meaning of neglecting velocity dispersion and pressure, and how these assumptions give rise to a pressureless fluid model.

We then examine the mathematical properties of the resulting equations, emphasizing their weakly hyperbolic nature. Unlike the Euler equations for gas dynamics, the pressureless dust equations admit degenerate eigenvalues and lack a restoring force associated with pressure. As a consequence, solutions can develop singular structures such as delta–waves, corresponding to the concentration of mass and momentum into infinitesimally thin regions. These features are closely related to the physical possibility of trajectory crossing and to the absence of an intrinsic mechanism preventing dust accumulation.

The second part of the section is devoted to the interaction between dust and gas via aerodynamic drag. We discuss how momentum exchange couples the two phases and introduces a characteristic stopping time that controls the degree of dynamical coupling. The drag force depends sensitively on both dust properties (grain size, internal density) and gas properties (density, temperature, mean free path), leading to different drag regimes such as the Epstein and Stokes limits. We highlight how the stopping time determines whether dust behaves as a tightly coupled tracer of the gas or as a dynamically independent component.

Finally, we review the main numerical strategies used to implement dust dynamics in hydrodynamical simulations. We contrast Eulerian and Lagrangian approaches and discuss how dust can be represented either as a collection of particles or as a continuum fluid. Particular attention is given to the strengths and limitations of each method, including their ability to handle multi–streaming, strong dust–gas coupling, and numerical artefacts

such as artificial clumping.

2.0.1 The Pressureless Fluid Model

The derivation of the governing equations for a dusty fluid closely follows the steps used to derive the standard Euler equations for a gas. It is useful to repeat this derivation here to explicitly show the assumptions and modeling choices that lead to the pressureless approximation.

From the Boltzmann Equation to Continuity

The dust fluid can be described by a particle distribution function $f(t, \vec{x}, \vec{v})$, such that

$$f(t, \vec{x}, \vec{v})\, d^3\vec{x}\, d^3\vec{v} \tag{2.1}$$

represents the number of particles with positions near $\vec{x}$ and velocities near $\vec{v}$ at time t. From this, the number density of dust is obtained by integrating over velocities:

$$n(\vec{x}, t) = \int f(t, \vec{x}, \vec{v})\, d^3\vec{v}. \tag{2.2}$$

Similarly, the mean of any quantity A over the distribution is

$$\langle A \rangle = \frac{1}{n} \int A f\, d^3\vec{v}. \tag{2.3}$$

The evolution of f is governed by the Boltzmann equation:

$$\frac{\partial f}{\partial t} + \vec{v} \cdot \nabla_{\vec{x}} f + \vec{F} \cdot \nabla_{\vec{v}} f = \left(\frac{\partial f}{\partial t} \right)_{\text{coll}}, \tag{2.4}$$

where $\vec{F}$ represents external forces, and $(\partial f/\partial t)_{\text{coll}}$ is the collision operator, describing changes in f due to particle interactions.
Multiplying Eq. (2.4) by the particle mass m and integrating over velocity gives

$$\int m \frac{\partial f}{\partial t} d^3\vec{v} + \int m v_j \frac{\partial f}{\partial x_j} d^3\vec{v} = \int m \left(\frac{\partial f}{\partial t} \right)_{\text{coll}} d^3\vec{v}, \tag{2.5}$$

where Einstein summation over repeated indices is implied.
The first term yields the rate of change of the mass density:

$$\partial_t \int m f d^3\vec{v} = \partial_t(\rho) \quad \text{with} \quad \rho = n\langle m \rangle. \tag{2.6}$$

The second term gives the divergence of the momentum density:

$$\partial_{x_j} \int m v_j f d^3\vec{v} = \partial_{x_j}(\rho u_j), \quad \rho u_j = \langle \rho v_j \rangle. \tag{2.7}$$

Regarding the collision term, in the usual derivation of the Euler equations for the gas, it is neglected because the collisions between the particles are treated as elastic. Elastic collisions conserve the total mass, momentum, and energy of the fluid, and thus the integral over the entire velocity space of the variations of f due to collisions is zero.
For dust we can consider two limiting cases:

1. **Collisionless dust:** If the number density is low and collisions are rare, we set $(\partial f/\partial t)_{\text{coll}} = 0$.

2. **Sticky collisions:** If particles collide and stick together, moving with the barycenter velocity, then the total mass and momentum are conserved. In this case, the integral of the collision operator over velocity space still vanishes.

Hence, for both cases, the continuity equation follows:

$$\partial_t \rho + \partial_{x_j}(\rho u_j) = 0. \tag{2.8}$$

Notice how the cancellation of the collision term is independent of the actual collision rule chosen, as long as it conserves mass and momentum.

Momentum Equation and Pressureless Approximation

To derive the momentum equation, we multiply Eq. (2.4) by mv_i and integrate:

$$\int mv_i \frac{\partial f}{\partial t} d^3\vec{v} + \int mv_i v_j \frac{\partial f}{\partial x_j} d^3\vec{v} = \int mv_i \left(\frac{\partial f}{\partial t}\right)_{\text{coll}} d^3\vec{v}. \tag{2.9}$$

Splitting the velocity into the bulk and random components,

$$v_i = u_i + (v_i - u_i), \tag{2.10}$$

we have

$$\partial_{x_j} \int mv_i v_j f \, d^3\vec{v} = \partial_{x_j}(\rho u_i u_j) + \partial_{x_j} P_{ij}, \tag{2.11}$$

where

$$P_{ij} = \rho\langle (v_i - u_i)(v_j - u_j)\rangle \tag{2.12}$$

is the pressure tensor representing the internal velocity dispersion.
For dust in astrophysical contexts, we often neglect P_{ij} for two main reasons:

1. Dust velocities are typically dominated by the bulk motion of the fluid, with little dispersion.

2. Inelastic collisions dissipate random velocities quickly.

Thus, the dust momentum equation reduces to the pressureless fluid form:

$$\partial_t(\rho u_i) + \partial_{x_j}(\rho u_i u_j) = 0. \tag{2.13}$$

Equations (2.8) and (2.13) form a closed system, in contrast to the Euler equations for a gas, which require an equation of state and an energy equation.

One-Dimensional Form and Hyperbolicity

In 1D, the system reads

$$\partial_t \rho + \partial_x(\rho u) = 0, \tag{2.14}$$
$$\partial_t(\rho u) + \partial_x(\rho u^2) = 0. \tag{2.15}$$

Defining the conserved variables and fluxes:

$$\vec{U} = \begin{pmatrix} \rho \\ \rho u \end{pmatrix}, \qquad \vec{F} = \begin{pmatrix} \rho u \\ \rho u^2 \end{pmatrix}, \tag{2.16}$$

the system can be compactly written as

$$\partial_t \vec{U} + \partial_x \vec{F} = 0. \tag{2.17}$$

A key property of the Euler equations for gas is strict hyperbolicity. A system is deemed strictly hyperbolic if the Jacobian of the flux has

- Real and distinct eigenvalues, corresponding to the wave speeds at which perturbations travel,
- A complete set of independent eigenvectors, describing the directions of propagation.

To study the hyperbolicity of the pressureless gas equation, we first write the Jacobian of the flux as

$$A(\vec{U}) = \frac{\partial \vec{F}}{\partial \vec{U}} = \begin{pmatrix} 0 & 1 \\ -u^2 & 2u \end{pmatrix}. \tag{2.18}$$

Its eigenvalues are found by solving the characteristic polynomial:

$$\det(A - \lambda I) = (\lambda - u)^2 = 0 \quad \Rightarrow \quad \lambda_1 = \lambda_2 = u. \tag{2.19}$$

Since the eigenvalues are not distinct, the system is only weakly hyperbolic, in contrast to the Euler equations for a gas, which are strictly hyperbolic. Physically, this degeneracy arises because the dust has zero pressure and therefore no sound waves propagate through it.

Delta waves

One consequence of the weak hyperbolicity of the pressureless gas equations is the development of delta waves. Consider the initial value problem

$$\rho(x,0) = \begin{cases} \rho_L & \text{if } x < 0 \\ \rho_R & \text{if } x > 0 \end{cases} \qquad u(x,0) = \begin{cases} u_L & \text{if } x < 0 \\ u_R & \text{if } x > 0 \end{cases} \tag{2.20}$$

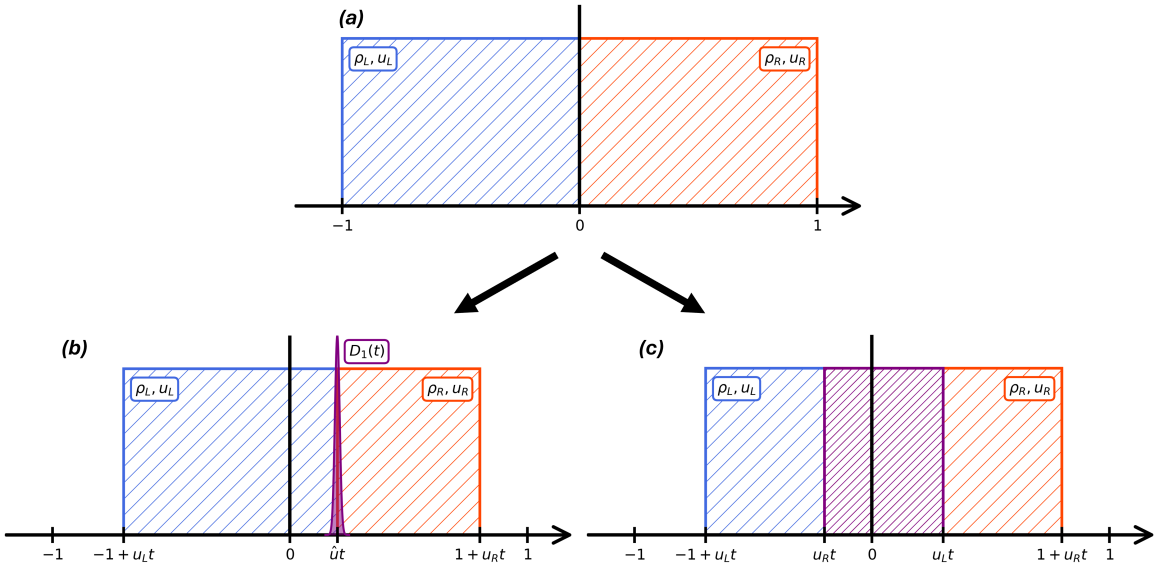

Figure 2.1: Schematic illustration of the initial value problem in Eq. 2.20 (a) and the two possible outcomes, depending on the chose microscopic collisional rule (b) and (c). (b) shows the development of the delta wave at the moving interface between the two fluids, arising from sticky collisions, while (c) shows the interpenetration of the two fluids, if no collisions are considered.

If $u_L < 0$ and $u_R > 0$, the two dust clouds will move apart and generate a vacuum state in between. Similarly, if $0 < u_L < u_R$ or $u_L < u_R < 0$, the two clouds will chase eachother but will not interact. The interesting behavior happens if $u_R < 0 < u_L$, when the two dust clouds move towards each other. In the sticky particles model, as soon as $t > 0$, dust grains will collect at the interface between the left and right states, building up a Dirac delta distribution. This dust clump will keep gathering dust and will start moving with speed $\hat{u}$. The formal solution to the pressureless fluid equations in this scenario is

$$\begin{cases} \rho(x,t) = \rho_L + (\rho_R - \rho_L)\theta(x - X(t)) + D_1(t)\delta(x - X(t)) \\ \rho(x,t)u(x,t) = \rho_L u_L + (\rho_R u_R - \rho_L u_L)\theta(x - X(t)) + D_2(t)\delta(x - X(t)), \end{cases} \tag{2.21}$$

where $X(t) = \hat{u}t$ is the location of the delta wave, θ is the Heaviside theta function, and $D_1(t)$ and $D_2(t)$ are the mass and momentum concentrated at the delta wave.
We can compare the density distributions for $t < 0$ and $t > 0$ and impose mass conservation to find $D_1(t)$ from

$$\rho_L + \rho_R = \rho_L[\hat{u}t - (-1 + u_L t)] + \rho_R[1 + u_R t - \hat{u}t] + D_1(t)D_1(t) = \rho_L t[u_L - \hat{u}] - \rho_R t[u_R - \hat{u}]. \tag{2.22}$$

Similarly, from momentum conservation, we can obtain

$$D_2(t) = \rho_L u_L t[u_L - \hat{u}] - \rho_R u_R t[u_R - \hat{u}]. \tag{2.23}$$

The speed of the delta wave can be described by the ratio between $D_2(t)$ and $D_1(t)$. By solving for $\hat{u}$ we find

$$\hat{u} = \frac{D_2(t)}{D_1(t)} = \frac{\sqrt{\rho_L}u_L + \sqrt{\rho_R}u_R}{\sqrt{\rho_L} + \sqrt{\rho_R}}. \tag{2.24}$$

Delta waves are thus the consequence of choosing a sticky particle microscopical collision rule. What would have happened if we had instead chosen the non-collisional dust model? In that case, no delta wave would develop. We would have an intersection region, $[-u_R t, u_L t]$, in which the density would be $\rho_L + \rho_R$. Still, the velocity would be multivalued, and thus cannot be modelled with a single dust fluid, which is defined by a single velocity. To properly model this situation, we would need to treat the leftward- and rightward-moving dust as two distinct fluids. For a single fluid, because the velocity field is multivalued, the structure of the pressureless equations breaks down, and we cannot recover a suitable solution to the initial value problem. However, in the next section, we show that a proper Riemann solver for dust dynamics remains possible even under the assumption of collisionless fluids. The intial value problem in Eq. 2.20 and the outcomes of the two collisional models described, are shown schematically in Fig. 2.1.

Riemann solvers for dust dynamics

One of the most straightforward numerical approaches to simulate hydrodynamical systems is the Godunov finite-volume method. We discretize the domain into N_c cells, each storing a vector of conserved quantities $\vec{U}_i$, with $i = 1, \ldots, N_\mathrm{c}$. Advancing the solution in time requires computing the intercell fluxes:

$$\vec{U}_i^{\,n+1} = \vec{U}_i^{\,n} + \frac{\Delta t}{\Delta x}\left(\vec{F}_{i-1/2} - \vec{F}_{i+1/2}\right), \tag{2.25}$$

where Δx is the cell width, Δt the timestep, and $t^{n+1} = t^n + \Delta t$.
Several Riemann solvers exist to compute the intercell fluxes. Some rely on strict hyperbolicity, such as the Exact Riemann Solver for the Euler equations, while others, such as the HLL solver, remain valid even for weakly hyperbolic systems. The latter is particularly convenient for dust dynamics, as it does not rely on the detailed wave structure of the Euler equations, but only on mass and momentum conservation.
However, it is also possible to construct dust–specific Riemann solvers that incorporate the chosen microscopic collision rule of dust particles. For example, Leveque (2004) developed a solver tailored for sticky particle dynamics, including the formation of delta waves. In this formulation, the intercell flux $F_{i-1/2}$ depends on the propagation speed $\hat{u}_{i-1/2}$ of the delta wave:

$$F_{i-1/2} = \begin{cases} F_{i-1}, & \hat{u}_{i-1/2} > 0, \\ \dfrac{1}{2}\left(F_{i-1} + F_i\right), & \hat{u}_{i-1/2} = 0, \\ F_i, & \hat{u}_{i-1/2} < 0. \end{cases} \tag{2.26}$$

with the additional constraint that $F_{i-1/2} = 0$ for diverging dust flows ($u_{i-1} < 0$ and $u_i > 0$).

The interface flux in this model is taken from the upwind side, with a symmetric average used when the delta wave is stationary.
An alternative dust Riemann solver is based on the collisionless particle rule. In this case, the left and right states evolve independently, as though the opposing side were absent. The flux contributions are:

$$F^{(L)}_{i-1/2} = \begin{cases} 0, & u_{i-1} < 0, \\ F_{i-1}, & u_{i-1} > 0, \end{cases} \qquad F^{(R)}_{i-1/2} = \begin{cases} F_i, & u_i < 0, \\ 0, & u_i > 0, \end{cases} \tag{2.27}$$

leading to the total flux:

$$F_{i-1/2} = F^{(L)}_{i-1/2} + F^{(R)}_{i-1/2} = \begin{cases} 0, & u_{i-1} < 0 \text{ and } u_i > 0, \\ F_{i-1}, & u_{i-1} > 0 \text{ and } u_i > 0, \\ F_i, & u_{i-1} < 0 \text{ and } u_i < 0, \\ F_{i-1} + F_i, & u_{i-1} > 0 \text{ and } u_i < 0. \end{cases} \tag{2.28}$$

Although this solver mimics collisionless behavior, it *cannot* truly represent multi-streaming or crossing trajectories in an Eulerian framework. Because the dust fluid model enforces a single-valued velocity field, two equal but opposite dust streams entering a central cell will numerically merge into a nearly zero net momentum, rather than passing through each other. This reflects the intrinsic limitation of the pressureless Euler system rather than of the solver itself.
To properly model crossing trajectories, one must adopt methods that allow multi-valued velocity distributions, such as:

- solving the full Boltzmann equation,
- using multiple dust fluids to represent distinct velocity streams,
- or employing (particle-based) methods.

These approaches retain information on the full phase-space structure that the pressureless Euler equations lose through trajectory crossing.

2.0.2 Aerodynamic drag

In the previous section, we examined the evolution equations for dust treated as a fluid in isolation. In realistic astrophysical environments, however, dust and gas coexist and

exchange momentum through aerodynamic drag. The coupled system is described by

$$\begin{cases} \partial_t \rho_\mathrm{g} + \nabla \cdot (\rho_\mathrm{g}\, \vec{v}_\mathrm{g}) = 0, \\ \partial_t \vec{v}_\mathrm{g} + (\vec{v}_\mathrm{g} \cdot \nabla)\vec{v}_\mathrm{g} = -\dfrac{\nabla P}{\rho_\mathrm{g}} + \dfrac{K}{\rho_\mathrm{g}}(\vec{v}_\mathrm{d} - \vec{v}_\mathrm{g}), \\ \partial_t \rho_\mathrm{d} + \nabla \cdot (\rho_\mathrm{d}\, \vec{v}_\mathrm{d}) = 0, \\ \partial_t \vec{v}_\mathrm{d} + (\vec{v}_\mathrm{d} \cdot \nabla)\vec{v}_\mathrm{d} = -\dfrac{K}{\rho_\mathrm{d}}(\vec{v}_\mathrm{d} - \vec{v}_\mathrm{g}). \end{cases} \tag{2.29}$$

The aerodynamic drag terms transfer momentum between dust and gas. Their strength is encoded in the drag coefficient K, which is commonly expressed in terms of the stopping time t_s as

$$K = \frac{\rho_\mathrm{d}\rho_\mathrm{g}}{\rho_\mathrm{d} + \rho_\mathrm{g}} \frac{1}{t_\mathrm{s}}. \tag{2.30}$$

To better understand the effect of drag, we define the velocity difference

$$\vec{\Delta v} \equiv \vec{v}_\mathrm{d} - \vec{v}_\mathrm{g}. \tag{2.31}$$

Keeping only the drag contributions to the evolution of $\vec{\Delta v}$ yields

$$\left(\frac{d\vec{\Delta v}}{dt}\right)_\mathrm{drag} = -\left(\frac{K}{\rho_\mathrm{d}} + \frac{K}{\rho_\mathrm{g}}\right)\vec{\Delta v} = -\frac{\vec{\Delta v}}{t_\mathrm{s}}. \tag{2.32}$$

The solution is a simple exponential decay:

$$\left[\vec{\Delta v}(t)\right]_\mathrm{drag} = \vec{\Delta v}(0)\, e^{-t/t_\mathrm{s}}. \tag{2.33}$$

Thus, dust–gas drift is damped on the stopping time: small grains (short t_s) quickly adapt to the gas velocity, while large grains preserve their inertia for longer.

Drag regimes

The stopping time depends on the properties of the grain and the surrounding gas. The relevant physical regime is determined by the Knudsen number,

$$\mathrm{Kn} = \frac{\lambda_\mathrm{mfp}}{a}, \tag{2.34}$$

where a is the grain radius and λ_mfp the gas mean free path.
If $\mathrm{Kn} > 4/9$, individual gas molecules collide with the grain more frequently than gas–gas collisions occur. The grain, therefore, interacts with molecular bombardment rather than a continuous flow: this is the Epstein regime (Epstein, 1924).
If $\mathrm{Kn} < 4/9$, the gas behaves as a fluid around the grain, leading to the Stokes regime. In that case, the drag depends on the particle Reynolds number,

$$\mathrm{Re}_\mathrm{p} = \frac{2a\, v_\mathrm{rel}}{\nu}, \tag{2.35}$$

where ν is the gas kinematic viscosity and $v_{\rm rel}$ is the dust–gas relative speed.
The aerodynamic drag force on a single grain can be written as

$$\vec{F}_{\rm D} = -\frac{1}{2} C_{\rm D}\, \pi a^2\, \rho_{\rm g}\, v_{\rm rel}\, \vec{v}_{\rm rel}, \tag{2.36}$$

where $C_{\rm D}$ is a dimensionless drag coefficient. Note that $v_{\rm rel}$ refers to the grain-level relative velocity, not the bulk fluid difference Δv.
In the Epstein limit,

$$C_{\rm D}^{\rm Ep} = \frac{8\,\langle v_{\rm th}\rangle}{3\,v_{\rm rel}}, \tag{2.37}$$

where $\langle v_{\rm th}\rangle$ is the mean thermal speed of the gas molecules. Substituting into Eq. (2.36), the stopping time defined as the momentum–loss timescale,

$$t_{\rm s} = \frac{m_{\rm d}\, v_{\rm rel}}{F_{\rm D}} = \frac{a\,\rho_{\rm gr}}{\langle v_{\rm th}\rangle\,\rho_{\rm g}}, \tag{2.38}$$

is independent of $v_{\rm rel}$ because $C_{\rm D} \propto v_{\rm rel}^{-1}$. Here $m_{\rm d} = (4\pi/3)a^3\rho_{\rm gr}$ is the grain mass.
In the Stokes regime, $C_{\rm D}$ depends on $\mathrm{Re_p}$. A standard parametrization, obtained from experimental data, is

$$C_{\rm D} = \begin{cases} 24\,\mathrm{Re_p}^{-1}, & \mathrm{Re_p} < 1, \\ 24\,\mathrm{Re_p}^{-0.6}, & 1 < \mathrm{Re_p} < 800, \\ 0.44, & \mathrm{Re_p} > 800, \end{cases} \tag{2.39}$$

following Probestein & Fassio (1970) (see also Cheng 2009 for a more recent fit). For $\mathrm{Re_p} < 1$, $C_{\rm D} \propto v_{\rm rel}^{-1}$ and the stopping time remains independent of $v_{\rm rel}$. At higher Reynolds numbers, however, this property no longer holds, and the stopping time becomes velocity-dependent due to nonlinear drag effects.

2.0.3 Methods for dust dynamics in hydrodynamical simulations

Simulating dust dynamics within hydrodynamical environments is a non-trivial task, primarily because dust behaves neither exactly as a pressureless fluid nor as a set of purely ballistic particles. Depending on the physical regime, dust may be strongly coupled to gas, weakly coupled, or even exhibit multiple simultaneous velocity streams. Numerical modeling must therefore address two separate decisions:

- whether the underlying simulation framework is Eulerian or Lagrangian, and

- whether dust itself is evolved Eulerianly (as a fluid) or Lagrangianly (as particles).

These choices give rise to four distinct approaches, each with strengths and limitations depending on the physical processes or scales of interest. Fig. 2.2 illustrates schematically the differences between these four approaches.

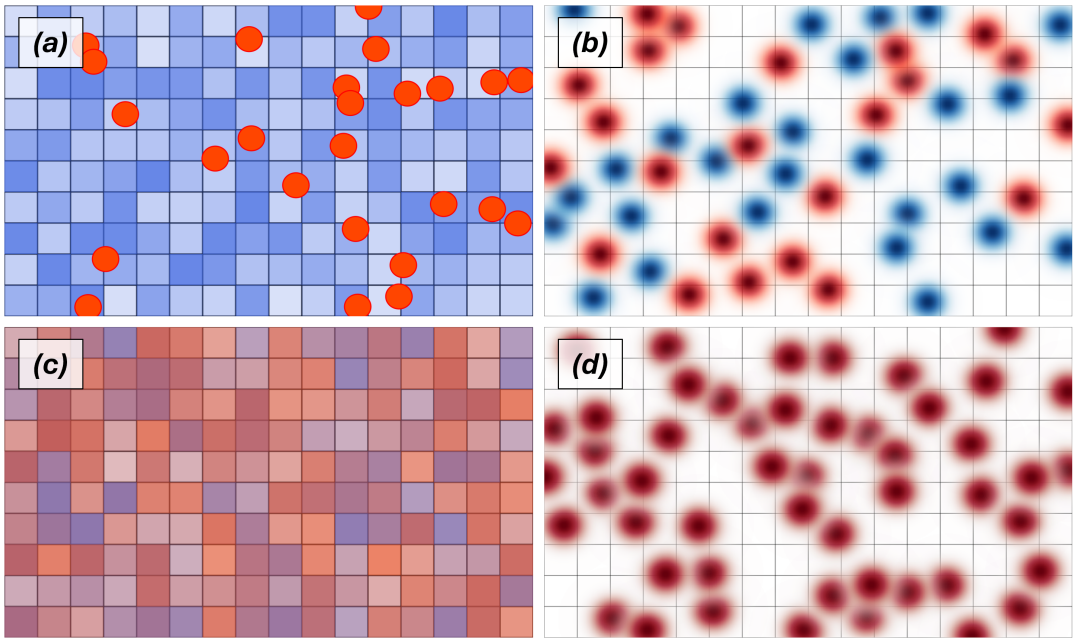

Figure 2.2: Schematic illustration of the four approaches to dust dynamics within hydrodynamical simulations. (*a*) Eulerian framework – Lagrangian dust, (*b*) Lagrangian framework – Lagrangian dust, (*c*) Eulerian framework – Eulerian dust, and (*d*) Lagrangian framework – Eulerian dust.

Eulerian framework – Lagrangian dust

In an Eulerian hydrodynamical code, dust can be represented as a collection of Lagrangian "superparticles," each evolved according to a system of ordinary differential equations:

$$\begin{cases} \dfrac{d\vec{x}_\mathrm{d}}{dt} = \vec{v}_\mathrm{d}, \\ \dfrac{d\vec{v}_\mathrm{d}}{dt} = \vec{a}_\mathrm{d} - \dfrac{\vec{v}_\mathrm{d} - \vec{v}_\mathrm{g}}{t_\mathrm{s}}. \end{cases} \tag{2.40}$$

where $\vec{a}_\mathrm{d}$ contains all accelerations other than aerodynamic drag (e.g., gravity, radiation pressure). Back-reaction of dust onto gas is commonly implemented through kernel-weighted deposition of drag forces onto the Eulerian grid (e.g., Yang & Johansen (2016)), allowing momentum exchange to remain conservative.
Because dust is treated as discrete particles, this method naturally handles multistreaming: particles can cross, diverge, or converge without any need for special treatment. This makes particle-based dust modeling well-suited for problems involving dust trapping, streaming instabilities, or the development of filamentary dust concentrations.
However, challenges arise near or below the grid resolution. Dust particles located deep inside a single Eulerian gas cell experience nearly identical drag forces and pressure gradients, which can prevent them from escaping local minima in the gas distribution. This

effect can lead to unphysical particle trapping or artificial clumping at cell centers ("grid-locking"; Price & Federrath (2010)). Additionally, because particle sampling introduces discrete noise, Lagrangian dust components typically display higher numerical noise than Eulerian dust fluids, especially in low-density regions.

Eulerian framework – Eulerian dust

In this configuration, dust is treated as a separate pressureless fluid co-evolved on the Eulerian grid alongside gas. Dust dynamics are advanced using a Riemann Solver, examples of dust-specific solvers were discussed in Section 2.0.1. This approach benefits from the stability, low noise, and mature numerical infrastructure of Eulerian finite-volume methods.
A major limitation, however, is that a single-valued Eulerian velocity field cannot represent multistreaming or crossing trajectories. Regions where dust should develop counter-flowing streams are instead forced into a single averaged velocity, artificially enhancing dust concentration and suppressing physically expected phase-space structure.
Because pressureless fluids lack a restoring force (no pressure gradients), the Eulerian dust equations are prone to catastrophic clumping. To maintain numerical stability, explicit or implicit numerical diffusion (or slope limiting) is generally required, though at the cost of damping small-scale structure.
Despite these issues, Eulerian–Eulerian schemes remain effective for modeling moderately or strongly coupled grains, large-scale transport, and global dust–gas momentum exchange.

Lagrangian framework – Lagrangian dust

Within a Lagrangian hydrodynamics framework such as SPH, dust can be treated as an additional population of Lagrangian particles. The Two-Fluid SPH method (Laibe & Price, 2012a) evolves gas and dust as separate particle sets, each with its own density estimate and smoothing length.
This approach mirrors the Eulerian–Lagrangian case by allowing multistreaming and by enabling dust to move independently of gas. It also preserves SPH's advantages, such as exact advection and automatic adaptivity.
However, SPH dust suffers from similar problems of particle trapping inside gas distributions when the spatial resolution of the gas is too low. More importantly, the two-fluid SPH formulation encounters severe temporal and spatial resolution constraints when the dust is tightly coupled to the gas. Laibe & Price (2012a) derived a spatial resolution requirement

$$\Delta x \lesssim c_s t_{\mathrm{s}}, \tag{2.41}$$

which becomes prohibitively restrictive for small grains with very short stopping times. In such cases, resolving drag forces requires extremely small timesteps and large numbers of particles, making two-fluid SPH computationally expensive.

Lagrangian framework – Eulerian dust

This last combination may seem contradictory at first and therefore requires some context. By Eulerian dust evolved in a Lagrangian framework, I do not mean that dust is evolved on a grid separate from the SPH particles describing the gas. Instead, dust is treated as a continuous fluid defined on the same discretization elements as the gas, without introducing a separate set of particles. In this sense, the evolution of dust resembles that of an Eulerian field, but defined on the moving, discretized SPH particle distribution.
In SPH codes, this approach is implemented via the so–called One–Fluid model, introduced by Laibe & Price (2014b). Each SPH particle represents a mixture of gas and dust, whose total density is given by

$$\rho = \rho_{\rm g} + \rho_{\rm d}. \tag{2.42}$$

The velocity used to advect the particles is the barycentric velocity of the mixture,

$$\vec{v} = \frac{\rho_{\rm g}\vec{v}_{\rm g} + \rho_{\rm d}\vec{v}_{\rm d}}{\rho_{\rm g} + \rho_{\rm d}}. \tag{2.43}$$

In addition, two auxiliary variables are introduced: the dust fraction ϵ and the dust–gas relative velocity $\vec{\Delta v}$,

$$\epsilon = \frac{\rho_{\rm d}}{\rho}, \qquad \vec{\Delta v} = \vec{v}_{\rm d} - \vec{v}_{\rm g}. \tag{2.44}$$

Under this change of variables and reference frame, the two–fluid system in Eq. 2.29 can be rewritten as

$$\begin{cases} \dfrac{d\rho}{dt} = -\rho\,(\nabla\cdot\vec{v}), \\ \dfrac{d\epsilon}{dt} = -\dfrac{1}{\rho}\,\nabla\cdot\left[\epsilon(1-\epsilon)\,\rho\,\vec{\Delta v}\right], \\ \dfrac{d\vec{v}}{dt} = -\dfrac{\nabla P}{\rho} - \dfrac{1}{\rho}\,\nabla\cdot\left[\epsilon(1-\epsilon)\,\rho\,\vec{\Delta v}\otimes\vec{\Delta v}\right], \\ \dfrac{d\vec{\Delta v}}{dt} = \dfrac{\nabla P}{\rho(1-\epsilon)} - \dfrac{\vec{\Delta v}}{t_{\rm s}} - (\vec{\Delta v}\cdot\nabla)\vec{v} + (\vec{\Delta v}\cdot\nabla\epsilon)\,\vec{\Delta v} + (2\epsilon-1)\,(\vec{\Delta v}\cdot\nabla)\vec{\Delta v}. \end{cases} \tag{2.45}$$

By discretizing these equations consistently within the SPH formalism, the mixture can be evolved self–consistently. Gas and dust variables can then be recovered by applying the inverse transformations

$$\begin{cases} \rho_{\rm g} = \rho(1-\epsilon), \\ \rho_{\rm d} = \rho\epsilon, \\ \vec{v}_{\rm g} = \vec{v} - \epsilon\,\vec{\Delta v}, \\ \vec{v}_{\rm d} = \vec{v} + (1-\epsilon)\,\vec{\Delta v}. \end{cases} \tag{2.46}$$

The One–Fluid model has been implemented in the PHANTOM code (Price et al., 2018), although currently only within the terminal velocity approximation (Youdin & Goodman,

2005). Under this approximation, the relative velocity is assumed to rapidly reach a steady–state value determined by the gas pressure gradient,

$$\vec{\Delta v} \simeq \frac{\nabla P}{\rho_{\rm g}} t_{\rm s}. \tag{2.47}$$

This approximation is valid for tightly coupled grains but breaks down for particles with stopping times comparable to or larger than the computational timestep. The restriction arises because, as in Eulerian–Eulerian models, the lack of multiple streaming and the absence of a restoring force can lead to artificial dust clumping onto individual SPH particles, violating the foundational SPH assumption that fluid quantities should be smoothed over multiple smoothing lengths.

3

Modeling dust dynamics in `OpenGadget3` I. SPH implementation of the One-Fluid model

Published as
Tedeschi-Prades et al., 2025, Astronomy & Astrophysics, Volume 704, n. A92

Full authors list:
Giovanni Tedeschi-Prades; Til Birnstiel; Klaus Dolag; Barbara Ercolano; Mark Hutchison.

Dust dynamics plays a critical role in astrophysical processes and has been modeled in hydrodynamical simulations using various approaches. Among particle-based methods like Smoothed Particle Hydrodynamics (SPH), the One-Fluid model has proven to be highly effective for simulating gas-dust mixtures.
This study presents the implementation of the One-Fluid model in `OpenGadget3`, introducing improvements to the original formulation. These enhancements include time-dependent artificial viscosity and conductivity, as well as a novel treatment of dust diffusion using a pressure-like term.
The improved model is tested using a suite of dust dynamics benchmark problems: DUSTYBOX, DUSTYWAVE, and DUSTYSHOCK, with the latter extended to multidimensional scenarios, as well as a dusty Sedov-Taylor blast wave. Additional tests include simulations of Cold Keplerian Disks, dusty protoplanetary disks, and Kelvin-Helmholtz instabilities to evaluate the model's robustness in more complex flows.
The implementation successfully passes all standard benchmark tests. It demonstrates stability and accuracy in both simple and complex simulations. The new diffusion term improves the handling of flows with large dust-to-gas ratios and low drag coefficients, although limitations of the One-Fluid model in these regimes remain.
The enhanced One-Fluid model is a reliable and robust tool for simulating dust dynamics in `OpenGadget3`. While it retains some limitations inherent to the original formulation, the introduced improvements expand its applicability and address some challenges in gas-dust dynamics.

3.1 Introduction

Dust, alongside gas, is one of the fundamental components of most astrophysical systems. From the cold interstellar medium to molecular clouds and protoplanetary disks, dust is observed ubiquitously across the Milky Way and other galaxies. Importantly, in many cases, dust is the only directly observable component, detected through its millimeter and infrared thermal emission. Gas properties are often inferred indirectly from dust observations, typically assuming a constant dust-to-gas ratio of 10^{-2}. Dust also plays a critical role in a variety of astrophysical processes. It contributes to the cooling of gas in star-forming regions (Vogelsberger et al., 2019), serves as a surface for chemical reactions that form complex molecules (Minissale et al., 2016), and provides the primary opacity source in many environments, impacting radiative transfer. Understanding the behavior of dust is therefore crucial not only for interpreting observations but also for modeling the physical conditions and evolution of astrophysical systems.

While dust is generally assumed to be perfectly coupled to the gas, this assumption breaks down either for larger dust-to-gas ratios or for larger grain sizes. Protoplanetary disks are the astrophysical environment where dust dynamics has been most thoroughly studied. Here, the uncoupling of dust and gas is essential to understand the evolution of the disk and the early stages of planet formation (Birnstiel, 2024). The radial drift of dust grains inward, one of the most important dynamical mechanisms to understand the evolution of disks, is caused by a slight difference between the rotational velocity of gas and of dust.

But disks are not the only objects in which the dynamics of dust grains play a role. The outflows of AGB stars are also believed to be accelerated by radiation pressure on dust grains, which consequently accelerates gas by aerodynamic drag. The same mechanism is also proposed to fuel outflows from "dusty torus" regions around AGNs (Soliman & Hopkins, 2023).

Given the importance of dust as a key component of many astrophysical systems, the inclusion of dust dynamics in modern hydrodynamical simulation codes has seen significant growth in recent years. One widely used approach involves post-processing gas-only simulations to study the behavior of dust. While this method improves upon the assumption of perfect coupling between dust and gas, it neglects the critical back-reaction of dust on gas motion. This omission can lead to inaccuracies, especially in scenarios with substantial dust-to-gas ratios or large grain sizes, where dust can significantly influence the gas dynamics. Self-consistent models of dust dynamics have been implemented in Eulerian codes, both as tracer particles, such as in `AREPO` (McKinnon et al., 2018a), or as additional fluids, like in `ATHENA++` (Huang & Bai, 2022) and `FARGO3D` (Benítez-Llambay et al., 2019). In Smoothed Particle Hydrodynamics (SPH) frameworks, the model that has achieved the most success in simulating dust dynamics is the One-Fluid model developed by (Laibe & Price, 2014b). This approach represents a significant departure from traditional multi-fluid methods by describing the motion of a gas-dust mixture self-consistently, rather than treating gas and dust as separate fluids coupled via aerodynamic drag. A notable implementation of the One-Fluid model, in the terminal velocity approximation limit, is in the

SPH code PHANTOM (Price et al., 2018). While this model has been largely successful at simulating the dynamics of small dust grains, it still presents challenges when simulating large dust grains. Most notably, (Laibe & Price, 2014c) observed that interpenetrating dust flows of weakly coupled grains can violate the assumptions underlying the fluid approximation for a pressureless dust model, causing the code to crash. More subtly, the inability of a pressureless fluid to capture the diffusive nature of a multi-valued velocity distribution leads to artificial clumping of dust, which can produce numerical failures or result in an overestimate of planetesimal formation timescales. This has led to the One-Fluid model to be used mostly in the terminal velocity approximation, where dust is assumed to couple efficiently to the gas on timescales smaller than the numerical integration timestep. One possible solution to this long-standing issue of simulating dust dynamics is the introduction of some form of dust pressure. For instance, Lynch et al. (2024) found that a treatment of dust pressure is needed in order to correctly simulate dust dynamics in debris disks.
In this paper, we will describe the implementation of the full (i.e., not in the terminal velocity approximation) One-Fluid model for dust dynamics in OpenGadget3. In particular, we will focus on improvements on the original implementation derived in (Laibe & Price, 2014c), such as time-dependent artificial viscosity and conductivity, and a model of dust diffusion via a dust pressure-like term to treat clumping in simulations of weakly coupled dust. While the limitations of the full One-Fluid model are still present in the most extreme cases, we show how these improvements can make the One-Fluid model suitable for a larger set of astrophysical applications compared to the terminal velocity approximation.
The paper is organized as follows. In Section 3.2 we show how the One-Fluid model is derived from the equations describing the motion of dust and gas separately and we show the discretization of the continuum equations in the SPH framework, including the improved viscosity and conductivity terms and the dust-pressure diffusion model. In Section 3.3, we show the performance of the code on a suite of standard numerical tests, as well as on more complex simulations such as the Cold Keplerian Disk and the Kelvin-Helmholtz instability. Finally, in Section 3.4, we present our conclusions and outline potential avenues for future development.

3.2 Methods

3.2.1 The One-Fluid model

Our implementation of dust dynamics relies and improves on the One-Fluid model introduced in Laibe & Price (2014b). The equations governing gas dynamics in the presence of dust closely resemble the standard Euler equations, with additional terms in the momentum and energy conservation equations to account for the drag exerted by dust on the gas

$$\partial_t \rho_g + \nabla \cdot (\rho_g \vec{v}_g) = 0, \tag{3.1}$$

$$\partial_t \vec{v}_g + (\vec{v}_g \cdot \nabla)\vec{v}_g = -\frac{\nabla P}{\rho_g} + \frac{K}{\rho_g}(\vec{v}_d - \vec{v}_g) + \vec{f}, \tag{3.2}$$

$$\partial_t u + (\vec{v}_g \cdot \nabla)u = -\frac{P}{\rho_g}(\nabla \cdot \vec{v}_g) + \frac{K}{\rho_g}(\vec{v}_d - \vec{v}_g)^2 + \Lambda, \tag{3.3}$$

where u and P are the internal energy and the pressure of the gas, respectively, and K is the drag coefficient determining how strongly the gas and the dust are coupled. For simplicity, we will ignore any external force acting on the fluid, $\vec{f}$, as well as any external heating, Λ.
The equations for dust are essentially describing a pressureless fluid, with an additional term in the momentum equation describing the drag of gas onto the dust, similar and opposite to the one in Eq.3.2

$$\partial_t \rho_d + \nabla \cdot (\rho_d \vec{v}_d) = 0, \tag{3.4}$$

$$\partial_t \vec{v}_d + (\vec{v}_d \cdot \nabla)\vec{v}_d = -\frac{K}{\rho_d}(\vec{v}_d - \vec{v}_g) + \vec{f}. \tag{3.5}$$

As in the case of the gas momentum equation, we will ignore the term $\vec{f}$ as well as any effect of gas pressure on the dust.
Laibe & Price (2014b) show that, by introducing a new and convenient set of variables, the equations describing the motion of gas and dust can be combined into a set of equations describing the motion of the mixture of both fluids. In particular, the following variables are used

$$\rho = \rho_g + \rho_d, \tag{3.6}$$

$$\epsilon = \frac{\rho_d}{\rho_g + \rho_d}, \tag{3.7}$$

$$\vec{v} = \frac{\rho_g \vec{v}_g + \rho_d \vec{v}_d}{\rho_g + \rho_d}, \tag{3.8}$$

$$\vec{\Delta v} = \vec{v_d} - \vec{v_g}, \tag{3.9}$$

which described, respectively, the total density (ρ), the fraction of dust density compared to the total density (ϵ), the barycenter velocity of the mixture ($\vec{v}$), and the velocity difference between the dust and the gas phases ($\vec{\Delta v}$). By using these variables, the Eq. 3.1 - 3.5 can be rewritten as

$$\frac{d\rho}{dt} = -\rho(\nabla \cdot \vec{v}), \tag{3.10}$$

$$\frac{d\epsilon}{dt} = -\frac{1}{\rho}\nabla \cdot [\epsilon(1-\epsilon)\rho\vec{\Delta v}], \tag{3.11}$$

$$\frac{d\vec{v}}{dt} = -\frac{\nabla P}{\rho} - \frac{1}{\rho}\nabla \cdot [\epsilon(1-\epsilon)\rho\vec{\Delta v}\vec{\Delta v}], \tag{3.12}$$

$$\frac{d\vec{\Delta v}}{dt} = \frac{\nabla P}{\rho(1-\epsilon)} - \frac{\vec{\Delta v}}{t_s} - (\vec{\Delta v} \cdot \nabla)\vec{v} + \vec{\Delta v}(\vec{\Delta v} \cdot \nabla)\epsilon + (2\epsilon - 1)(\vec{\Delta v} \cdot \nabla)\vec{\Delta v}, \tag{3.13}$$

$$\frac{du}{dt} = -\frac{P}{\rho(1-\epsilon)}\nabla \cdot (\vec{v} - \epsilon\vec{\Delta v}) + \epsilon(\vec{\Delta v} \cdot \nabla)u + \epsilon\frac{\vec{\Delta v}^2}{t_s}, \tag{3.14}$$

where we already introduced the Lagrangian, or co-moving, temporal derivative

$$\frac{d}{dt} = \frac{\partial}{\partial t} + (\vec{v} \cdot \nabla), \tag{3.15}$$

and the stopping time t_s, which is related to the drag coefficient by

$$t_s = \frac{\rho\epsilon(1-\epsilon)}{K}. \tag{3.16}$$

Notice that Eq. 3.13 is slightly different from the form presented in Laibe & Price (2014b) due to an error in the original derivation.
Although Eq. 3.13 is the most compact form for the evolution of $\vec{\Delta v}$, here we present two alternative formulations. The first formulation is the one derived in Lebreuilly et al. (2019)

$$\begin{aligned}\frac{d\vec{\Delta v}}{dt} = \frac{\nabla P}{\rho(1-\epsilon)} - \frac{\vec{\Delta v}}{t_s} - (\vec{\Delta v}\cdot\nabla)\vec{v} + \frac{1}{2}\nabla[(2\epsilon-1)\vec{\Delta v}\cdot\vec{\Delta v}] \\ + (1-\epsilon)\vec{\Delta v}\times[\nabla\times(1-\epsilon)\vec{\Delta v}] - \epsilon\vec{\Delta v}\times[\nabla\times\epsilon\vec{\Delta v}].\end{aligned} \tag{3.17}$$

The second version will be the one that we will use for the later SPH discretization

$$\begin{aligned}\frac{d\vec{\Delta v}}{dt} = \frac{\nabla P}{\rho(1-\epsilon)} - \frac{\vec{\Delta v}}{t_s} - (\vec{\Delta v}\cdot\nabla)\vec{v} + \frac{1}{2}\nabla[(2\epsilon-1)\vec{\Delta v}\cdot\vec{\Delta v}] \\ + \vec{\Delta v}(\vec{\Delta v}\cdot\nabla)\epsilon + (2\epsilon-1)(\vec{\Delta v}\cdot\nabla)\vec{\Delta v} - |\vec{\Delta v}|^2\nabla\epsilon - \frac{1}{2}(2\epsilon-1)\nabla|\vec{\Delta v}|^2.\end{aligned} \tag{3.18}$$

While all these three formulations seem very different, they are actually equivalent to one another, and each may be useful under particular circumstances.
We note here that Eq. 3.17 and Eq. 3.18 share a common property: under scalar product with $\vec{\Delta v}$ itself, the second row of both equations cancels out. It is easier to see in Eq. 3.17, where all terms in the second row are the result of cross-products with $\vec{\Delta v}$, and thus perpendicular to it, but the same holds for Eq. 3.18 as well. This property will be crucial in deriving a suitable SPH discretization of Eq. 3.18.

3.2.2 SPH implementation of the One-Fluid model

In order to computationally solve the One-Fluid model presented above, we need to find a suitable discretization that conserves mass, momentum, and energy. We will closely follow the derivation of Laibe & Price (2014c) (hereafter LP14b) for the well-known Smoothed Particle Hydrodynamics (SPH) framework, but corrected for the new terms in Eq. 3.18. We will also add time-dependent versions of the artificial viscosity and conductivity, a new dust diffusion term, and the explicit form of the equations for a parametrization of the dust fraction.

Density estimate

In SPH, the fluid is modeled as a collection of particles, each representing a portion of its mass. Any hydrodynamical quantity at a point in space $\vec{r}$ is then represented by a weighted average of the contribution of all neighboring particles, weighted by a kernel $\mathcal{W}(\vec{r}_a, h_a)$, depending on the distance $\vec{r}_a$ between the particle a and the point $\vec{r}$ and the smoothing length h_a of the particle. The kernel should satisfy several properties: compact support, continuity, radial symmetry, and the following normalization constraint

$$\int_V \mathcal{W}(|\vec{r}' - \vec{r}_a|, h_a) dV' = 1. \tag{3.19}$$

Some common choices for SPH kernels include the cubic (Monaghan & Lattanzio, 1985) and quintic spline (Morris, 1996) and the Wedland C2/C4 and C6 kernels (Wendland, 1995; Dehnen & Aly, 2012).
The density of a given particle is then determined by a weighted sum over the particle's neighbors

$$\rho(\vec{r}_a) = \sum_{b \in NGB} m_b \mathcal{W}(|\vec{r}_a - \vec{r}_b|, h_a), \tag{3.20}$$

where m_b is the mass of the b particle. Eq. 3.20 can be shown to be a solution to the continuity equation Eq. 3.10. The variable smoothing length h_a of each particle is computed to keep the mass inside the kernel volume of each particle fixed. In three dimensions, h_a is computed to fulfill the equation

$$\frac{4}{3}\pi h_a^3 \rho_a = N_{\text{ngb}} m_a, \tag{3.21}$$

where N_{ngb} is the chosen number of neighbors.

Differential operators

To obtain a suitable discretization of the One-Fluid model equations we will employ the following SPH differential operators (Price, 2012)

$$\nabla \cdot \vec{A} \underset{\text{SPH}}{=} \rho_a \sum_b m_b \left[\frac{\vec{A}_a \cdot \nabla \mathcal{W}_{ab}(h_a)}{\Omega_a \rho_a^2} + \frac{\vec{A}_b \cdot \nabla \mathcal{W}_{ab}(h_b)}{\Omega_b \rho_b^2} \right], \tag{3.22}$$

$$(\vec{A} \cdot \nabla) f \underset{\text{SPH}}{=} \frac{1}{\rho_a \Omega_a} \sum_b m_b (f_b - f_a) \vec{A}_a \cdot \nabla \mathcal{W}_{ab}(h_a), \tag{3.23}$$

$$(\vec{A} \cdot \nabla) \vec{B} \underset{\text{SPH}}{=} \frac{1}{\rho_a \Omega_a} \sum_b m_b (\vec{B}_b - \vec{B}_a) \vec{A}_a \cdot \nabla \mathcal{W}_{ab}(h_a), \tag{3.24}$$

where $\mathcal{W}_{ab}(h_a) = \mathcal{W}(|\vec{r}_a - \vec{r}_b|, h_a)$, $\Omega_a = \left(1 + \frac{h_a}{\rho_a d} \frac{\partial \rho_a}{\partial h_a}\right)$, with d the number of dimensions, and $\vec{A}$, $\vec{B}$ and f are arbitrary vector and scalar fields.
We will apply these operators and, step by step, ensure that mass, total momentum, and total energy are conserved.

Mass conservation

The first requirement that the SPH discretization should meet is the conservation of mass. The total mass of the mixture is conserved by construction by the SPH framework, and since there is no dust-mass evolution mechanism taken into account in the One-Fluid model (e.g., sputtering, erosion, or astration), the individual dust and gas masses should also be conserved. Dust mass conservation requires that

$$\frac{dM_d}{dt} \underset{\text{SPH}}{=} \sum_a m_a \frac{d\epsilon_a}{dt} = 0. \tag{3.25}$$

By applying the divergence operator Eq. 3.22 to Eq. 3.11 we obtain the following SPH discretization

$$\begin{aligned}\frac{d\epsilon_a}{dt} = -\sum_b m_b \Bigg[&\frac{\epsilon_a(1-\epsilon_a)}{\Omega_a \rho_a} \vec{\Delta v}_a \cdot \nabla_a \mathcal{W}_{ab}(h_a) \\ &+ \frac{\epsilon_b(1-\epsilon_b)}{\Omega_b \rho_b} \vec{\Delta v}_b \cdot \nabla_a \mathcal{W}_{ab}(h_b) \Bigg].\end{aligned} \tag{3.26}$$

Using the anti-symmetry of the kernel gradient $\nabla \mathcal{W}_{ab}(h_a) = -\nabla \mathcal{W}_{ab}(h_b)$, it is straightforward to show that the total mass of dust is conserved.

Momentum conservation

While dust and gas can exchange momentum, the total momentum of the mixture should be conserved. A suitable SPH discretization of Eq. 3.12 is constrained by

$$\sum_a m_a \frac{d\vec{v}_a}{dt} = 0. \tag{3.27}$$

Compared to the gas-only momentum conservation equation, the One-Fluid model version, Eq. 3.12, has a first term that represents the usual pressure gradient acceleration and a new term specific to the One-Fluid model. The pressure gradient term can be solved as usual in SPH

$$\left(\frac{d\vec{v}_a}{dt}\right)_{\nabla P} = -\sum_b m_b \left[\frac{P_a}{\Omega_a \rho_a^2} \nabla \mathcal{W}_{ab}(h_a) + \frac{P_b}{\Omega_b \rho_b^2} \nabla \mathcal{W}_{ab}(h_b) \right]. \tag{3.28}$$

The Eq. 3.18 also presents a similar term, which we discretize in the same manner

$$\begin{aligned}\left(\frac{d\vec{\Delta v}_a}{dt}\right)_{\nabla P} &= -\frac{1}{1-\epsilon_a}\left(\frac{d\vec{v}_a}{dt}\right)_{\nabla P} = \\ &= \frac{1}{1-\epsilon_a} \sum_b m_b \left[\frac{P_a}{\Omega_a \rho_a^2} \nabla \mathcal{W}_{ab}(h_a) + \frac{P_b}{\Omega_b \rho_b^2} \nabla \mathcal{W}_{ab}(h_b) \right].\end{aligned} \tag{3.29}$$

Back to the momentum conservation, the second term of Eq. 3.12 can be discretized by applying the differential operator Eq. 3.22

$$\left(\frac{d\vec{v}_a}{dt}\right)_{OFM} = -\sum_b m_b \Bigg[\frac{\epsilon_a(1-\epsilon_a)\vec{\Delta v}_a}{\Omega_a\rho_a}\vec{\Delta v}_a \cdot \nabla\mathcal{W}_{ab}(h_a) + \\ + \frac{\epsilon_b(1-\epsilon_b)\vec{\Delta v}_b}{\Omega_b\rho_b}\vec{\Delta v}_b \cdot \nabla\mathcal{W}_{ab}(h_b)\Bigg]. \tag{3.30}$$

Both Eq. 3.28 and Eq. 3.30 conserve momentum as required by Eq. 3.27 and can thus be used as SPH discretizations of the corresponding continuum equations.

Energy conservation

The mixture of dust and gas should also conserve its total energy, and from this requirement, we will derive the SPH discretizations for both the internal energy u and the remaining terms in the $\vec{\Delta v}$ evolution.
The total energy of the mixture is the sum of the total energy of the gas phase (kinetic and internal energy) and the total energy of the dust phase (which only has a kinetic component, as it is modeled as a pressureless fluid)

$$\begin{aligned} E &= \int_V \left(\frac{1}{2}m_g\vec{v}_g^2 + m_g u + \frac{1}{2}m_d\vec{v}_d^2\right)dV = \\ &= \int_V m\left(\frac{1}{2}\vec{v}^2 + (1-\epsilon)u + \frac{1}{2}\epsilon(1-\epsilon)\vec{\Delta v}^{\,2}\right)dV \underset{\text{SPH}}{=} \\ &\underset{\text{SPH}}{=} \sum_a m_a\left(\frac{1}{2}\vec{v}_a^2 + (1-\epsilon_a)u_a + \frac{1}{2}\epsilon_a(1-\epsilon_a)\vec{\Delta v}_a^{\,2}\right), \end{aligned} \tag{3.31}$$

where in the second row we introduced the One-Fluid model variables and in the final row we discretized the integral as a sum over all particles.
By computing the time derivative of Eq. 3.31 we obtain the constraint that the SPH discretizations for Eq. 3.18 and 3.14 should satisfy

$$\frac{dE}{dt} = \sum_a m_a \Bigg[\vec{v}_a \cdot \frac{d\vec{v}_a}{dt} + \epsilon_a(1-\epsilon_a)\vec{\Delta v}_a \cdot \frac{d\vec{\Delta v}_a}{dt} + \\ \left((1-2\epsilon_a)\frac{\vec{\Delta v}_a^{\,2}}{2} - u_a\right)\frac{d\epsilon_a}{dt} + (1-\epsilon_a)\frac{du_a}{dt}\Bigg] = 0. \tag{3.32}$$

The first term we want to discretize is the PdV term in Eq. 3.14. By combining all terms in which pressure is present, we obtain the following constraint from energy conservation

$$\begin{aligned}
&\sum_a m_a(1-\epsilon_a)\left(\frac{du_a}{dt}\right)_{PdV} = \\
&\quad = -\sum_a m_a\left[\vec{v}_a\cdot\left(\frac{d\vec{v}_a}{dt}\right)_{\nabla P} + \epsilon_a(1-\epsilon_a)\vec{\Delta v}_a\cdot\left(\frac{d\vec{\Delta v}_a}{dt}\right)_{\nabla P}\right] = \\
&\quad = -\sum_a m_a\left[(\vec{v}_a-\epsilon_a\vec{\Delta v}_a)\cdot\left(\frac{d\vec{v}_a}{dt}\right)_{\nabla P}\right] \\
&\quad = \sum_a m_a\sum_b m_b\frac{P_a}{\Omega_a\rho_a^2}(\vec{v}_a-\epsilon_a\vec{\Delta v}_a)\cdot\nabla\mathcal{W}_{ab}(h_a)+ \\
&\qquad + \sum_a m_a\sum_b m_b\frac{P_b}{\Omega_b\rho_b^2}(\vec{v}_a-\epsilon_a\vec{\Delta v}_a)\cdot\nabla\mathcal{W}_{ab}(h_b),
\end{aligned} \tag{3.33}$$

where in the last step we have substituted $\left(\frac{d\vec{v}_a}{dt}\right)_{\nabla P}$ from Eq. 3.28. Swapping the summation indices of the second term and using the anti-symmetry of the kernel, we obtain

$$\left(\frac{du_a}{dt}\right)_{PdV} = \frac{P_a}{\Omega_a\rho_a^2(1-\epsilon_a)}\sum_b m_b[(\vec{v}_a-\epsilon_a\vec{\Delta v}_a-\vec{v}_b+\epsilon_b\vec{\Delta v}_b)\cdot\nabla W_{ab}(h_a)]. \tag{3.34}$$

The second term in Eq. 3.14 is also constrained from energy conservation. By comparing the terms in Eq. 3.32 involving the internal energy

$$\sum_a m_a(1-\epsilon_a)\left(\frac{du_a}{dt}\right)_{(\vec{\Delta v}\cdot\nabla)u} = \sum_a m_a u_a\frac{d\epsilon_a}{dt}. \tag{3.35}$$

By substituting Eq. 3.26 and swapping indices in the double summation, we obtain

$$\left(\frac{du_a}{dt}\right)_{(\vec{\Delta v}\cdot\nabla)u} = \frac{\epsilon_a}{\Omega_a\rho_a}\sum_b m_b(u_b-u_a)\vec{\Delta v}_a\cdot\nabla\mathcal{W}_{ab}(h_a). \tag{3.36}$$

We will now deal with Eq. 3.18, describing the evolution of $\vec{\Delta v}$. The first term, similar to the pressure gradient force term, was already dealt with and its SPH discretization is described in Eq. 3.29. The second term in Eq. 3.18 relating to the drag will be integrated implicitly and addressed later, together with the last term in Eq. 3.14. The SPH discretization for the third term in Eq. 3.18 can be derived from energy conservation, by equating the first two terms in Eq. 3.32:

$$\sum_a m_a\epsilon_a(1-\epsilon_a)\vec{\Delta v}_a\cdot\left(\frac{d\vec{\Delta v}_a}{dt}\right)_{-(\vec{\Delta v}\cdot\nabla)\vec{v}} = -\sum_a m_a\vec{v}_a\cdot\left(\frac{d\vec{v}_a}{dt}\right)_{OFM}. \tag{3.37}$$

Also in this case, we simply substitute Eq. 3.30 and swap summation indices to obtain

$$\left(\frac{d\vec{\Delta v}_a}{dt}\right)_{-(\vec{\Delta v}\cdot\nabla)\vec{v}} = -\frac{1}{\Omega_a\rho_a}\sum_b m_b[(\vec{v}_b - \vec{v}_a)\vec{\Delta v}_a \cdot \nabla\mathcal{W}_{ab}(h_a)]. \tag{3.38}$$

The fourth term in Eq. 3.18 is also constrained by energy conservation, but comparing now the second and third terms in Eq. 3.32

$$\sum_a m_a\epsilon_a(1-\epsilon_a)\vec{\Delta v}_a \cdot \left(\frac{d\vec{\Delta v}_a}{dt}\right)_4 = -\sum_a m_a(1-2\epsilon_a)\frac{\vec{\Delta v}_a^{\,2}}{2}\frac{d\epsilon_a}{dt}. \tag{3.39}$$

By substituting Eq. 3.26 and swapping summation indices we obtain

$$\begin{aligned}\left(\frac{d\vec{\Delta v}_a}{dt}\right)_4 = \frac{1}{2\Omega_a\rho_a}\sum_b m_b[&(1-2\epsilon_a)\vec{\Delta v}_a^{\,2} - \\ &- (1-2\epsilon_b)\vec{\Delta v}_b^{\,2}]\nabla\mathcal{W}_{ab}(h_a).\end{aligned} \tag{3.40}$$

It is crucial to note here that the remaining terms in Eq. 3.18 cannot be constrained by energy conservation, as they cancel out when multiplied scalarly by $\vec{\Delta v}$, as noted when they were first introduced (see Section 3.2.1). Hence, their SPH discretization will be obtained directly by use of SPH differential operators (Eq. 3.22 - 3.24), and we will only later check that energy is conserved. We propose the following SPH discretizations

$$\vec{\Delta v}(\vec{\Delta v}\cdot\nabla)\epsilon \underset{\text{SPH}}{=} \frac{\vec{\Delta v}_a}{\Omega_a\rho_a}\sum_b m_b(\epsilon_b-\epsilon_a)\vec{\Delta v}_a\cdot\nabla\mathcal{W}_{ab}(h_a), \tag{3.41}$$

$$(2\epsilon-1)(\vec{\Delta v}\cdot\nabla)\vec{\Delta v} \underset{\text{SPH}}{=} \frac{2\epsilon_a-1}{\Omega_a\rho_a}\sum_b m_b(\vec{\Delta v}_b-\vec{\Delta v}_a)\vec{\Delta v}_a\cdot\nabla\mathcal{W}_{ab}(h_a), \tag{3.42}$$

$$-|\vec{\Delta v}|^2\nabla\epsilon \underset{\text{SPH}}{=} -\frac{|\vec{\Delta v}_a|^2}{\Omega_a\rho_a}\sum_b m_b(\epsilon_b-\epsilon_a)\nabla\mathcal{W}_{ab}(h_a), \tag{3.43}$$

$$-\frac{1}{2}(2\epsilon-1)\nabla|\vec{\Delta v}|^2 \underset{\text{SPH}}{=} -\frac{2\epsilon_a-1}{\Omega_a\rho_a}\sum_b m_b(\vec{\Delta v}_b-\vec{\Delta v}_a)\cdot\vec{\Delta v}_a\nabla\mathcal{W}_{ab}(h_a). \tag{3.44}$$

It can be shown that these SPH discretizations conserve energy if inserted into Eq. 3.32. One quantity that this set of SPH discretizations fails to conserve is the dust momentum in the absence of drag forces. However, deviations from conservation are only of order $\mathcal{O}(\vec{\Delta v}^{\,2})$, and thus negligible for strong drag. In the weak drag case, we also observed little deviation from conservation in the benchmark tests we show in Section 3. In general, the conservation of dust momentum for zero drag is a property that has still not been enforced by previous SPH discretizations of the One-Fluid model, and could be an area of future improvement.

3.2.3 Timestepping

`OpenGadget3` employs, for integrating in time all hydrodynamical variables, a leapfrog scheme in a Kick-Drift-Kick (KDK) form. This scheme ensures second-order accuracy for all variables and is implemented as in Springel (2005).
We indicate with $\vec{W} = (\epsilon, \vec{v}, \vec{\Delta v}, u)$ the vector of primitive quantities of the One-Fluid model, except for density, which is computed by the kernel density estimate in Eq. 3.20. The $\vec{W}$ vector and the positions of the particles, $\vec{x}$, are evolved using predicted quantities to compute the increments of the second kick

$$\begin{aligned}
\vec{W}_{n+1/2} &= \vec{W}_n + \frac{1}{2}\Delta t \left(\frac{d\vec{W}}{dt}\right)_n, \\
\vec{W}^*_{n+1} &= \vec{W}_n + \Delta t \left(\frac{d\vec{W}}{dt}\right)_n, \\
\vec{x}_{n+1} &= \vec{x}_n + \Delta t\ \vec{v}_{n+1/2}, \\
\vec{W}_{n+1} &= \vec{W}_{n+1/2} + \frac{1}{2}\Delta t \left(\frac{d\vec{W}}{dt}\right)^*_{n+1}.
\end{aligned} \tag{3.45}$$

3.2.4 Implicit integration of the drag term

The last two terms in Eq. 3.18 and Eq. 3.14 can be numerically integrated both explicitly and implicitly[1]. As shown in LP14b, explicit integration is subject to the timestep requirement $\Delta t < t_s$, which becomes prohibitive for strong drag. Implicit integration, instead, does not require such a timestep requirement and can be implemented via operator splitting.
We show here the procedure for the first kick in the KDK scheme, but the same is applied also to compute the predicted quantities during the drift and for the second kick. By separating the drag term from all the other contributions, we can rewrite Eq. 3.18 and 3.14 as

$$\begin{cases}
\dfrac{d\vec{\Delta v}}{dt} = \left(\dfrac{d\vec{\Delta v}}{dt}\right)_0 - \dfrac{\vec{\Delta v}}{t_s}, \\
\dfrac{du}{dt} = \left(\dfrac{du}{dt}\right)_0 + \epsilon \dfrac{\vec{\Delta v}^{\,2}}{t_s}.
\end{cases} \tag{3.46}$$

[1] Technically, the integrator presented here should be referred to as "analytic", rather than "implicit". However, in practice, all integrators of drag that are not explicit are usually referred to as "implicit" in the literature.

To use operator splitting, we first solve the corresponding homogeneous system by ignoring the drag terms and using the SPH discretization derived above

$$\begin{cases} \overline{\vec{\Delta}v}_{n+1/2} = \vec{\Delta}v_n + \frac{1}{2}\Delta t \left(\frac{d\vec{\Delta}v}{dt}\right)_{n,0}, \\ \overline{u}_{n+1/2} = u_n + \frac{1}{2}\Delta t \left(\frac{du}{dt}\right)_{n,0}. \end{cases} \tag{3.47}$$

The Ordinary Differential Equation corresponding to the system in Eq. 3.46 can be solved analytically as

$$\begin{cases} \frac{d\vec{\Delta}v}{dt} = -\frac{\vec{\Delta}v}{t_s}, \\ \frac{du}{dt} = \epsilon\frac{\vec{\Delta}v^2}{t_s}, \end{cases} \quad \rightarrow \quad \begin{cases} \vec{\Delta}v(t) = \vec{\Delta}v_0 e^{-\frac{t}{t_s}}, \\ u(t) = u_0 - \frac{1}{2}\epsilon\vec{\Delta}v_0^2\left[e^{-\frac{2t}{t_s}} - 1\right]. \end{cases}$$

By discretizing this solution to half a timestep as in the first kick, and by using the solutions to the homogeneous system as initial conditions, we obtain the final result

$$\begin{cases} \vec{\Delta}v_{n+1/2} = \left[\vec{\Delta}v_n + \frac{1}{2}\Delta t\left(\frac{d\vec{\Delta}v}{dt}\right)_{n,0}\right] e^{-\frac{\Delta t}{2t_s}}, \\ u_{n+1/2} = u_n + \frac{1}{2}\Delta t\left(\frac{du}{dt}\right)_{n,0} - \frac{1}{2}\epsilon_{n+1/2}\overline{\vec{\Delta}v}^2_{n+1/2}\left[e^{-\frac{\Delta t}{t_s}} - 1\right]. \end{cases} \tag{3.48}$$

3.2.5 Time-dependent Artificial Viscosity

In order to dampen post-shock oscillations, prevent particle interpenetration, and treat contact discontinuities, SPH frameworks often employ artificial viscosity mechanisms (Price, 2008). Our artificial viscosity implementation follows the steps of LP14b, with some modifications and improvements that make the implementation time-dependent. We show both the conservative (C) and the non-conservative (NC) approaches to artificial viscosity, as shown in LP14b, an artificial dissipation term for $\vec{\Delta}v$, and a simple unifying scheme to combine both conservative and non-conservative viscosities.

Conservative Artificial Viscosity

A simple momentum-conserving formulation for artificial viscosity can be written as

$$\left(\frac{d\vec{v}}{dt}\right)^C_{AV} = \sum_b m_b(1 - \overline{\epsilon}_{ab})\nu_{ab}\overline{\nabla\mathcal{W}}_{ab}, \tag{3.49}$$

where the overline stands for the mean value between particle a and particle b of a given quantity. For instance, $\overline{\epsilon}_{ab} = \frac{1}{2}(\epsilon_a + \epsilon_b)$.

A possible choice for the viscosity ν_{ab} is the one employed by LP14b

$$\nu_{ab} = \frac{v_{ab}^{sig}}{\overline{\rho}_{ab}} \vec{v}_{ab}^{g} \cdot \hat{\vec{r}}_{ab}, \tag{3.50}$$

where $\vec{v}_{ab}^{g} = \vec{v}_{a}^{g} - \vec{v}_{b}^{g}$ is the difference between the gas velocities of the particles, $\hat{\vec{r}}_{ab}$ is the unit vector pointing from particle a to particle b and v_{ab}^{sig} is the signal velocity.
The signal velocity is used to switch on and off the artificial viscosity depending on whether the gas in the particles is approaching or distancing

$$v_{ab}^{sig} = \begin{cases} \frac{1}{2}(c_{s,a} + c_{s,b} - \beta \vec{v}_{ab}^{g} \cdot \hat{\vec{r}}_{ab}) & \vec{v}_{ab}^{g} \cdot \hat{\vec{r}}_{ab} \leq 0, \\ 0 & \vec{v}_{ab}^{g} \cdot \hat{\vec{r}}_{ab} > 0, \end{cases} \tag{3.51}$$

where c_s is the sound speed of the particle and $\beta = 3$ following Beck et al. (2016).
In our implementation, we adopt two additional terms in the formulation of viscosity

$$\nu_{ab} = \frac{v_{ab}^{sig}}{\overline{\rho}_{ab}} \alpha_{ab} \vec{v}_{ab}^{g} \cdot \hat{\vec{r}}_{ab}, \tag{3.52}$$

where α_{ab} is the symmetrized viscosity coefficient, and is used to reduce artificial viscosity when it is not needed, i.e., away from shocks. This framework was already implemented (Dolag et al., 2005; Beck et al., 2016) and extensively tested in `OpenGadget3` by Marin-Gilabert et al. (2022), we only added the precaution of using the gas velocity instead of the mixture velocity when computing α_{ab}.
In order to make this formulation conservative, we need to add matching terms in the equations of $\vec{\Delta v}$ and u. In particular, for $\vec{\Delta v}$ we simply have

$$\left(\frac{d\vec{\Delta v}_a}{dt}\right)_{AV}^{C} = \frac{1}{1-\epsilon_a} \sum_b m_b (1 - \overline{\epsilon}_{ab}) \nu_{ab} \overline{\nabla \mathcal{W}}_{ab}, \tag{3.53}$$

while the term in the internal energy equation is constrained from the equivalence already shown in Eq. 3.33

$$\sum_a m_a (1-\epsilon_a) \left(\frac{du_a}{dt}\right)_{AV}^{C} = -\sum_a m_a \left[\vec{v}_a^{g} \cdot \left(\frac{d\vec{v}_a}{dt}\right)_{AV}^{C}\right], \tag{3.54}$$

which gives, after inserting Eq. 3.49, expanding $\overline{\nabla \mathcal{W}_{ab}}$ and swapping summation indices

$$\left(\frac{du_a}{dt}\right)_{AV}^{C} = -\frac{1}{2(1-\epsilon_a)} \sum_b m_b (1 - \overline{\epsilon}_{ab}) \nu_{ab} \vec{v}_{ab}^{g} \cdot \hat{\vec{r}}_{ab} \overline{F}_{ab}, \tag{3.55}$$

where the kernel gradient has been expanded into direction and magnitude: $\nabla \mathcal{W}_{ab}(h_a) = \hat{\vec{r}}_{ab} F_{ab}(h_a)$.
As described by LP14b, only using Eq. 3.49, 3.53 and 3.55 leaves some oscillations in $\vec{\Delta v}$

in post-shock regions, which can be treated by using an additional dissipative term in the $\vec{\Delta}v$ equation. In particular, a momentum and energy-conserving formulation is

$$\left(\frac{d\vec{\Delta}v_a}{dt}\right)_{AD} = \frac{1}{\epsilon_a(1-\epsilon_a)}\sum_b \frac{v_{sig}^{\vec{\Delta}v}}{\overline{\rho}_{ab}}\overline{\epsilon}_{ab}(1-\overline{\epsilon}_{ab})\vec{\Delta}v_{ab}\cdot\hat{\vec{r}}_{ab}\overline{\nabla\mathcal{W}_{ab}}. \tag{3.56}$$

In order to conserve energy, a matching term in the internal energy equation must be added

$$\left(\frac{du_a}{dt}\right)_{AD} = -\frac{1}{2(1-\epsilon_a)}\sum_b m_b\frac{v_{sig}^{\Delta v}}{\overline{\rho}_{ab}}\overline{\epsilon}_{ab}(1-\overline{\epsilon}_{ab})(\vec{\Delta}v_{ab}\cdot\hat{\vec{r}}_{ab})^2\overline{F}_{ab}. \tag{3.57}$$

The signal velocity used here is different from the one used in the artificial viscosity. In particular, LP14b propose the following definition

$$v_{ab}^{sig,\Delta v} = \begin{cases} -\alpha_{\Delta v}\vec{\Delta}v_{ab}\cdot\hat{\vec{r}}_{ab} & \vec{\Delta}v_{ab}\cdot\hat{\vec{r}}_{ab}\leq 0, \\ 0 & \vec{\Delta}v_{ab}\cdot\hat{\vec{r}}_{ab} > 0, \end{cases} \tag{3.58}$$

with $\alpha_{\Delta v} = 1$. This concludes the conservative formulation for artificial viscosity. As already found by LP14b, this formulation works well for low $\vec{\Delta}v$, but fails in shocks with large $\vec{\Delta}v$, as can be seen in the DUSTYSHOCK benchmark test, at low drag. For this reason, LP14b introduced an alternative, non-conservative, formulation of artificial viscosity that shows a much better performance on the DUSTYSHOCK test.

Nonconservative Artificial Viscosity

The nonconservative formulation of artificial viscosity we employed is very similar to the one derived by LP14b, with the only improvement of making it adaptive and time-dependent by using the viscosity in Eq. 3.52

$$\left(\frac{d\vec{v}_a}{dt}\right)_{AV}^{NC} = \sum_b m_b\nu_{ab}\overline{\nabla\mathcal{W}}_{ab}, \tag{3.59}$$

$$\left(\frac{d\vec{\Delta}v_a}{dt}\right)_{AV}^{NC} = \frac{1}{1-\epsilon_a}\sum_b \frac{v_{ab}^{sig}}{\overline{\rho}_{ab}}\vec{\Delta}v_{ab}\cdot\hat{\vec{r}}_{ab}\overline{\nabla\mathcal{W}_{ab}}, \tag{3.60}$$

where the signal velocity employed in the second equation is the one in Eq. 3.51.
The viscous heating term is also added as in Eq. 3.55, slightly modified to adjust to the different term in the $\frac{d\vec{v}}{dt}$ equation

$$\left(\frac{du_a}{dt}\right)_{AV}^{NC} = -\frac{1}{2(1-\epsilon_a)}\sum_b m_b\nu_{ab}\vec{v}_{ab}^{\,g}\cdot\hat{\vec{r}}_{ab}\overline{F}_{ab}(h_a). \tag{3.61}$$

Unified framework

In order to derive a unified framework that could be used under all circumstances, instead of relying on two options to be left to the user, we use the shock indicator R_a to interpolate between the two formulations. The shock indicator is defined as in Cullen & Dehnen (2010)

$$R_a = \frac{1}{\rho_a} \sum_b \text{sign}(\nabla \cdot \vec{v}_b) m_b \mathcal{W}_{ab}. \tag{3.62}$$

For non-shocked regions, $R_a \simeq 0$, while for shocked regions $R_a \simeq \pm 1$. This indicator is already used to compute the symmetrized viscosity coefficient α_{ab}, to determine the strength of the shock.
Starting from the shock indicator, we define a new variable, σ_{ab}, to interpolate between the conservative and the non-conservative formulations of artificial viscosity. σ_{ab} is defined as a sigmoid function

$$\sigma_{ab} = \frac{1}{1 + e^{\alpha(\overline{|R|}_{ab} - \beta)}}, \tag{3.63}$$

where appropriate values for α and β where found to be $\alpha = 20$ and $\beta = 0.25$. This makes the sigmoid function steep enough to quickly transition from one formulation to the other. In particular, for shocked regions $\sigma_{ab} \simeq 0$, while for non-shocked regions $\sigma_{ab} \simeq 1$.
The interpolation between conservative and non-conservative formulations is performed as follows

$$\left(\frac{d\vec{v}}{dt}\right)_{AV} = \left(\frac{d\vec{v}}{dt}\right)^{C}_{AV} \sigma_{ab} + \left(\frac{d\vec{v}}{dt}\right)^{NC}_{AV} (1 - \sigma_{ab}), \tag{3.64}$$

$$\left(\frac{d\Delta\vec{v}}{dt}\right)_{AV} = \left[\left(\frac{d\Delta\vec{v}}{dt}\right)^{C}_{AV} + \left(\frac{d\Delta\vec{v}}{dt}\right)_{AD}\right] \sigma_{ab} + \left(\frac{d\Delta\vec{v}}{dt}\right)^{NC}_{AV} (1 - \sigma_{ab}), \tag{3.65}$$

$$\left(\frac{du}{dt}\right)_{AV} = \left[\left(\frac{du}{dt}\right)^{C}_{AV} + \left(\frac{du}{dt}\right)_{AD}\right] \sigma_{ab} + \left(\frac{du}{dt}\right)^{NC}_{AV} (1 - \sigma_{ab}). \tag{3.66}$$

This is the framework that will be used throughout all the benchmark tests and simulations.

3.2.6 Time-dependent Artificial Conductivity

Artificial conductivity is introduced to treat discontinuities in internal energy (Price, 2008). Here we employ its time-dependent formulation introduced in Beck et al. (2016), which is also used in `OpenGadget3` (see Marin-Gilabert et al., 2022; Groth et al., 2023), adapted for the One-Fluid model.,

$$\left(\frac{du_a}{dt}\right)_{AC} = \frac{1}{1 - \epsilon_a} \sum_b \frac{m_b}{\overline{\rho}_{ab}} (u_b - u_a) \alpha^C_{ab} v^{sig,c}_{ab} \overline{F}_{ab}, \tag{3.67}$$

where the signal velocity for the artificial conductivity is defined as

$$v^{sig,c}_{ab} = \sqrt{\frac{|P_a - P_b|}{\overline{\rho(1-\epsilon)}_{ab}}}, \tag{3.68}$$

and the conductivity coefficient α_a^C is made time-dependent by computing it from the internal energy and its gradient:

$$\alpha_a^C = \frac{h_a}{2}\frac{|\nabla u_a|}{|u_a|}. \tag{3.69}$$

3.2.7 Dust fraction parametrization

In order to ensure that the dust fraction, ϵ, remains bounded between 0 and 1, a suitable parametrization is needed. Here we follow (Ballabio et al., 2018) and implement the following parametrization

$$\epsilon = \frac{s^2}{1+s^2} \rightarrow s = \sqrt{\frac{\epsilon}{1-\epsilon}}. \tag{3.70}$$

The corresponding continuum equation for the variable s can be obtained by substituting the parametrization in Eq. 3.11 and by using that

$$\frac{ds}{dt} = \frac{(1+s^2)^2}{2s}\frac{d\epsilon}{dt}, \tag{3.71}$$

which results in

$$\frac{ds}{dt} = -\frac{(1+s^2)^2}{2\rho s}\nabla\cdot\left[\frac{s^2}{(1+s^2)^2}\,\rho\vec{\Delta v}\right]. \tag{3.72}$$

Since the s at the denominator could cause numerical issues if a given SPH particle has very little or no dust, we split the $s^2 = s\cdot s$ inside the divergence operator to obtain the following

$$\frac{ds}{dt} = -\frac{(1+s^2)^2}{2\rho}\nabla\cdot\left[\frac{s}{(1+s^2)^2}\,\rho\vec{\Delta v}\right] - \frac{1}{2}(\vec{\Delta v}\cdot\nabla)s. \tag{3.73}$$

This last formulation is thus discretized using the SPH operators Eq. 3.22-3.23 and rearranged to obtain

$$\begin{aligned}\frac{ds_a}{dt} = -\frac{1}{2}\sum_b m_b s_b\Bigg[&\frac{1}{\Omega_a\rho_a}\vec{\Delta v}_a\cdot\nabla\mathcal{W}_{ab}(h_a)\\ &+\frac{(1+s_a^2)^2}{\Omega_b\rho_b(1+s_b^2)^2}\vec{\Delta v}_b\cdot\nabla\mathcal{W}_{ab}(h_b)\Bigg].\end{aligned} \tag{3.74}$$

The other equations and corresponding SPH discretizations do not need to be recomputed; we simply replace ϵ when needed with the parametrization in Eq. 3.70, in terms of the new variable s.

3.2.8 Artificial Diffusion of dust

One of the consequences of modeling dust as a pressureless fluid is its tendency, when the drag coefficient is sufficiently low, to form clumps of dust that cannot self-regulate and dissipate. This leads to large gradients in ϵ, which will lead, eventually, to code

crashes. In order to alleviate this tendency, we tried including some form of dust diffusion in the One Fluid model. Our first attempt was to tackle directly the evolution of ρ_d by adding a Gradient Diffusion flux to its continuity equation. While this may certainly be a valid approach in Eulerian simulations, we found this approach not suitable for our SPH framework, as it would require adding a term in the continuity equation for the total mixture density, and hence a mass transfer term between SPH particles.
We thus resorted to a different approach, the inclusion of a "pressure-like" term in the dust momentum conservation equation, as described by Klahr & Schreiber (2021) and Binkert (2023)

$$\frac{\partial}{\partial t}(\rho_d \vec{v}_d) + \nabla \cdot (\rho_d \vec{v}_d \vec{v}_d) = -\nabla\left(\frac{1}{3}\frac{D}{t_s}\rho_d\right) - \frac{\rho_d}{t_s}(\vec{v}_d - \vec{v}_g). \tag{3.75}$$

Translating this approach in the One Fluid model equations, we obtain the following continuum equations

$$\left(\frac{d\vec{v}}{dt}\right)_{diff} = -\frac{DK}{3\rho(1-\epsilon)^2}\nabla\epsilon = -\frac{2DKs}{3\rho}\nabla s, \tag{3.76}$$

$$\left(\frac{d\vec{\Delta v}}{dt}\right)_{diff} = -\frac{DK}{3\rho\epsilon(1-\epsilon)^2}\nabla\epsilon = -\frac{2sDK}{3\rho\epsilon}\nabla s, \tag{3.77}$$

where we already provide the formulation for the dust fraction parametrization (Eq. 3.70) that we will be using throughout the paper.
We propose the following suitable (momentum- and energy-conserving) SPH discretization for the previous equations

$$\left(\frac{d\vec{v}_a}{dt}\right)_{diff} = -\sum_b m_b \frac{2\overline{D}_{ab}\overline{s}_{ab}K}{3\rho_a\rho_b}(s_b - s_a)\overline{\nabla\mathcal{W}_{ab}}, \tag{3.78}$$

$$\left(\frac{d\vec{\Delta v}_a}{dt}\right)_{diff} = -\sum_b m_b \frac{2\overline{D}_{ab}\overline{s}_{ab}K}{3\rho_a\rho_b\overline{\epsilon}_{ab}}(s_b - s_a)\overline{\nabla\mathcal{W}_{ab}}, \tag{3.79}$$

$$\left(\frac{du_a}{dt}\right)_{diff} = \sum_b m_b \frac{2\overline{D}_{ab}\overline{s}_{ab}K}{3\rho_a\rho_b(1-\overline{\epsilon}_{ab})}(s_b - s_a)(\vec{v}_a^d - \vec{v}_b^d)\cdot\hat{r}_{ab}\overline{F}_{ab}, \tag{3.80}$$

where the additional heating term was added to conserve energy, similarly to how Eq. 3.55 is introduced.
We only need to specify what diffusion coefficient D to use. While some implementations of dust diffusion leave this choice to the user (as in Huang & Bai (2022)), we want our implementation to be adaptive and only trigger when needed in order to make the dust density field smoother in case of large gradients. We tried different approaches, but what was found to have the best results in terms of stability was the following. We define a new variable, χ_a, as the ratio between the dust fraction of particle a and the mean of the dust fractions of all neighboring particles

$$\chi_a = \frac{\epsilon_a}{\frac{1}{N_{ngb}}\sum_b \epsilon_b}. \tag{3.81}$$

We want diffusion to act on particles for which χ_a departs from unity. A suitable definition for D was found to be

$$D_a = \frac{(\chi_a - 1)^2}{\chi_a}. \tag{3.82}$$

In the SPH discretizations we use the symmetrical $\overline{D_{ab}}$ to ensure momentum and energy conservation.
This implementation could also be extended to describe a physical (or non-artificial) diffusion. This would require modeling the diffusion coefficient D based on gas or dust properties, rather than on the numerical prescription Eq. 3.82.
In principle, defining a spatially varying diffusion coefficient could allow shocks to develop, which would then require a dedicated artificial-viscosity term. To test whether this actually occurs, we examined the ratio between the compression and advection timescales as a shock indicator. These timescales are defined as

$$\tau_{\rm compr} = \frac{1}{|\nabla \cdot \vec{v}_{\rm d}|}, \quad \tau_{\rm adv} = \frac{h}{|\vec{v}_{\rm d}|}. \tag{3.83}$$

In a shock, compression is much faster than advection alone would allow, and thus $\tau_{\rm compr} < \tau_{\rm adv}$. We thus define the shock indicator for dust

$$C = \frac{h \cdot |\nabla \cdot \vec{v}_{\rm d}|}{|\vec{v}_{\rm d}|}, \tag{3.84}$$

so that a shock happens in the dust when $C > 1$.
In Fig. 3.1 we show the histograms of the values for C for the Cold Keplerian Disk and the Kelvin-Helmholtz instability simulations discussed, respectively, in Section 3.3.5 and Section 3.3.7. For all the drag coefficients K, and dust-to-gas ratios δ, in simulations the vast majority of particles show values of $C < 1$, indicating the absence of dust shocks.

3.3 Tests

In this section, we aim at showing both the accuracy and the stability of the implementation of the One-Fluid model in `OpenGadget3` shown in the previous section. We first show that the code is able to reproduce the expected results from standard dust dynamics tests: the DUSTYBOX, DUSTYWAVE, and DUSTYSHOCK test,s as well as a dusty Sedov-Taylor blast wave. For all these tests, we will not use our new dust diffusion model. Afterward, we test the stability of the code under more complex flows for various dust-to-gas ratios and drag coefficients. One of the major challenges in simulating dust dynamics is the tendency in low drag and high dust-to-gas ratio environments for dust to clump and create large gradients in dust fraction. To study these behaviors and find numerical solutions to make the code stable under these extreme circumstances, we simulate a Cold Keplerian Disk, a dusty protoplanetary disk, and a Kelvin-Helmholtz instability. These more complex flows

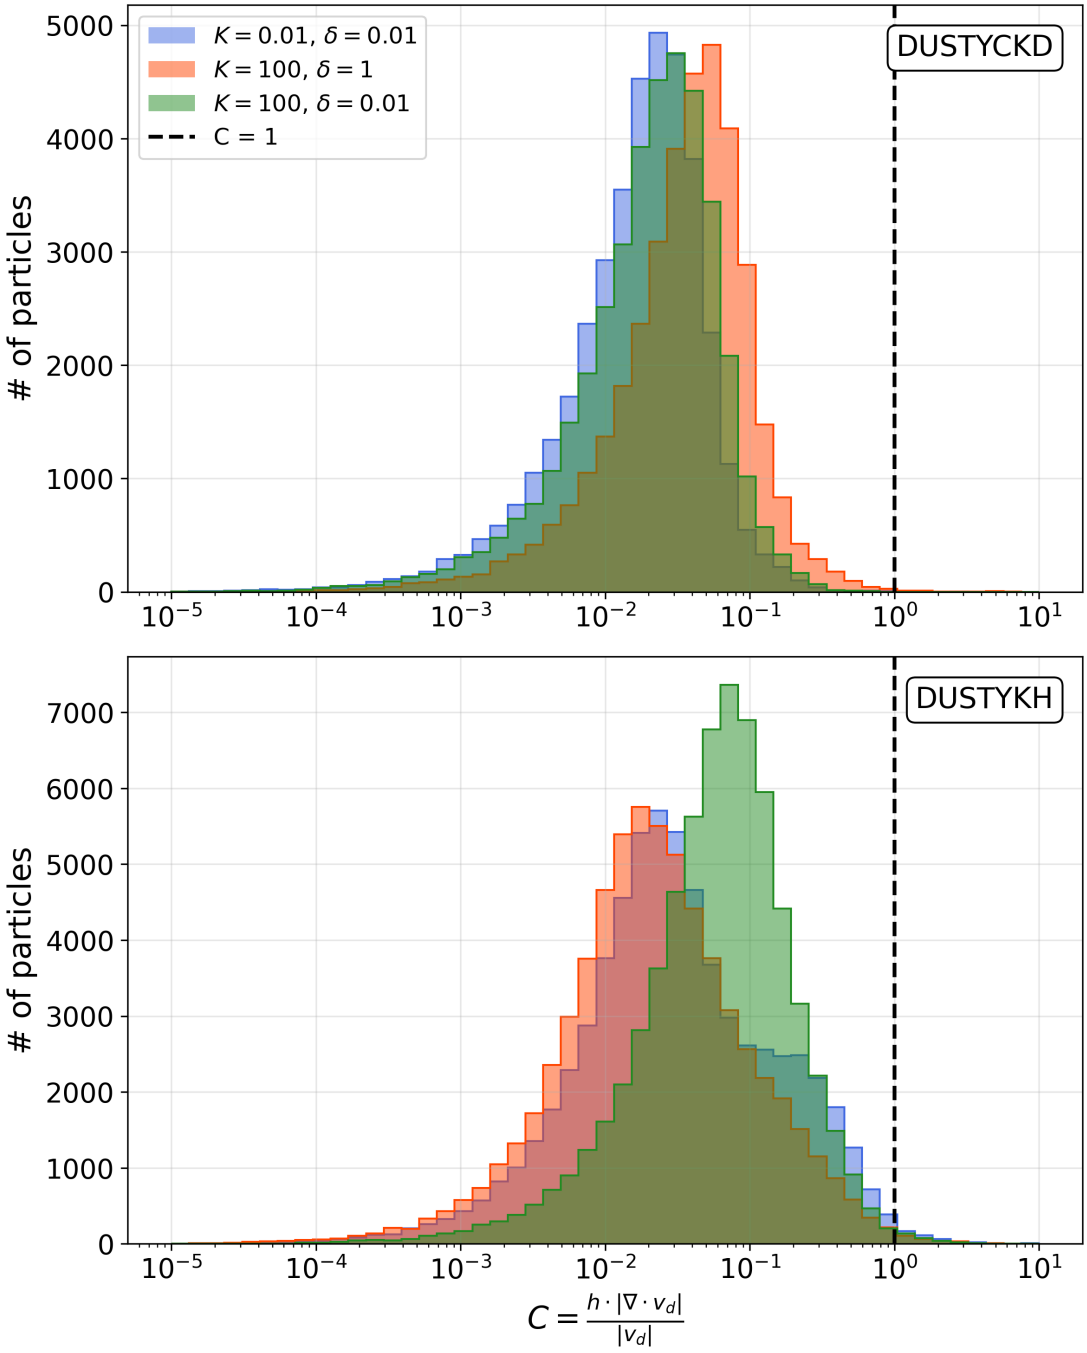

Figure 3.1: Distribution of the shock indicator defined in Eq. 3.84 for the Cold Keplerian Disk (top panel) and the Kelvin-Helmholtz instability (bottom panel) simulations.

will be simulated with dust diffusion turned on. All the values reported in the following section are in internal code units.

3.3.1 DUSTYBOX

The first benchmark test we perform is the so-called DUSTYBOX test. It consists of a mixture of gas and dust moving with respect to each other, but globally at rest. This is primarily used to test the integration of the drag term, since all other terms in the evolution equations are zero. The setup is extremely simple: 100 particles are equally spaced in one dimension, the density is set to $\rho_0 = 1$, and the dust fraction is set to $\epsilon_0 = 1/2$, meaning that there is an equal amount of dust and gas. The barycenter velocity is zero, $\vec{v} = 0$, so the mixture is at rest and the particles do not move in space. Additionally, the pressure is constant and set to $P_0 = 1$.

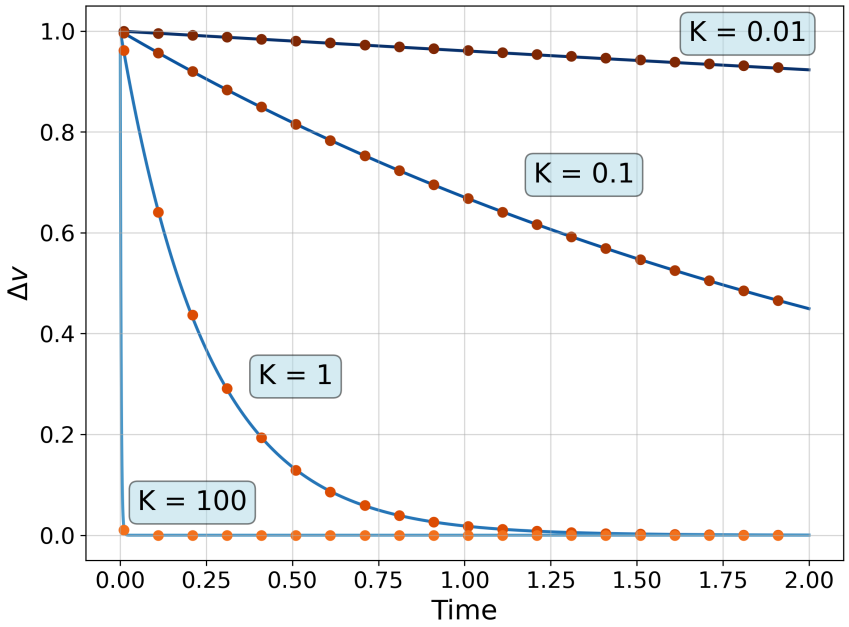

Figure 3.2: Result of the DUSTYBOX benchmark test. The continuous lines represent the analytical solution, while the dots are the simulation outputs. The results are shown for four different values of the drag coefficient K.

The only dynamics arises from the dust-gas velocity difference, which is initialized to $\vec{\Delta v}_0 = 1$. The expected behavior is for this initial velocity difference to decay exponentially, following the simple analytical solution

$$\vec{\Delta v}(t) = \vec{\Delta v}_0 e^{-\frac{t}{t_s}}. \tag{3.85}$$

The stronger the drag coefficient K, the smaller the stopping timescale t_s and the faster the exponential decay.

In Fig. 3.2, we show the results of our simulations for four different values of the drag coefficient. The simulation outputs show very good agreement with the expected behaviors. More quantitatively, we found the L1 norm to be always smaller than $< 0.1\%$. In addition, energy is also very well conserved: as the simulation evolves, the kinetic energy stored in the second term of the expression for the total energy in Eq. 3.31 is converted into internal energy of the gas, via drag heating. The resulting energy is conserved up to differences of 10^{-7}.

3.3.2 DUSTYWAVE

The next benchmark test is a simple linear wave of a mixture of dust and gas. Since dust is a pressureless fluid, the motion is primarily driven by gas pressure, which makes the gas wave move across the domain. The dust is dragged by the gas according to the strength of the drag coefficient. The wave is initialized by perturbing an equilibrium particle distribution.

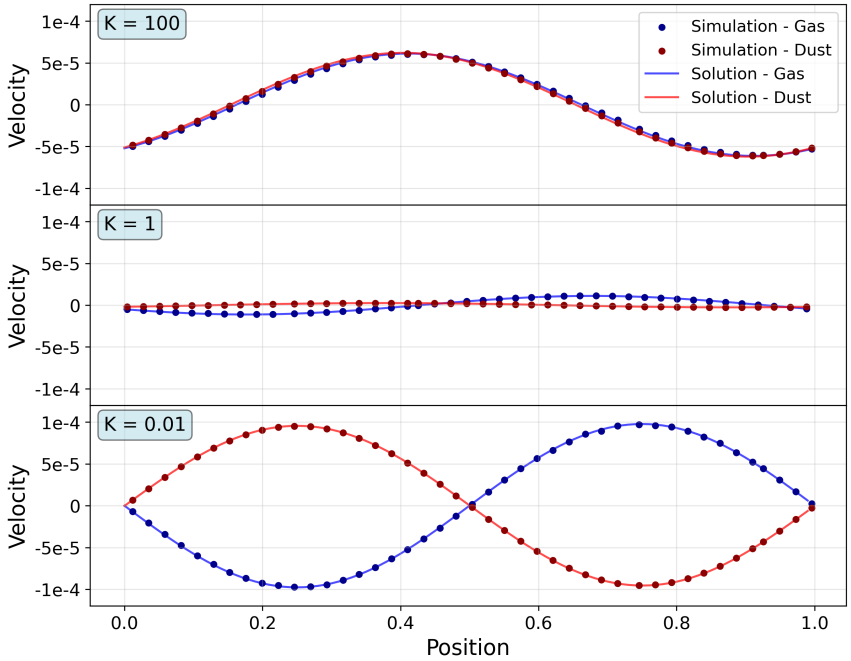

Figure 3.3: Results from the DUSTYWAVE benchmark test, for three different values of drag coefficient K. The continuous lines represent the analytical solution and the dots the output of the simulations. Gas and dust velocities are plotted separately.

The particles span a periodic domain $x \in (0, 1)$. The initial conditions for all variables are

$$\begin{aligned}
\rho(x, 0) &= \rho_0 + A \sin(2\pi x), \\
\vec{v}(x, 0) &= A \sin(2\pi x), \\
\epsilon(x, 0) &= 1/2, \\
\vec{\Delta v}(x, 0) &= 0,
\end{aligned}$$

where $A = 10^{-4}$ is the amplitude of the perturbation and $\rho_0 = 1$ is the background density value. The pressure is initialized with an isothermal equation of state, $P = \rho c_s^2$, with sound speed $c_s = 1$.
What makes the DUSTYWAVE a useful benchmark test is the availability of an analytical solution, derived by Laibe & Price (2011).
The results of our simulations are shown in Fig. 3.3, for three values of drag coefficient. Depending on the strength of the drag coefficient, the dynamics of dust and gas can differ widely. We find very good agreement between simulations and analytical solutions, with L_1 norms of the order of 10^{-7}. Also in this case energy is very well conserved. This is also achieved by artificial viscosity being active only in its conservative formulation, since the shock indicator here is very close to 0.

3.3.3 DUSTYSHOCK

The last suite of benchmark tests consists of a set of shocks. This is useful to test fully the implementation of the One-Fluid model, as shock environments require all of our artificial viscosity and conductivity terms. We explore two opposite scenarios: first, a weak drag regime, in which dust is expected to remain still as the shock propagates in the gas; secondly, a strong drag regime, where dust and gas move tightly coupled, and the shock propagates as a gas-only shock with modified sound speed. For both these regimes, we can compute a numerical solution using a Riemann Solver (here we chose an Exact Riemann Solver).
The DUSTYSHOCK test was already used in LP14b to test the One-Fluid model implementation in one dimension, but here we extend the test also to two dimensions.

Weak coupling

The gas is initialized as in a standard Sod shock tube test (Sod, 1978). The left and right states are

$$\begin{aligned}
\rho_L^g = 1 \quad \rho_R^g = 0.125, \\
\vec{v}_L^g = 0 \quad \vec{v}_R^g = 0, \\
P_L = 1 \quad P_R = 0.1.
\end{aligned}$$

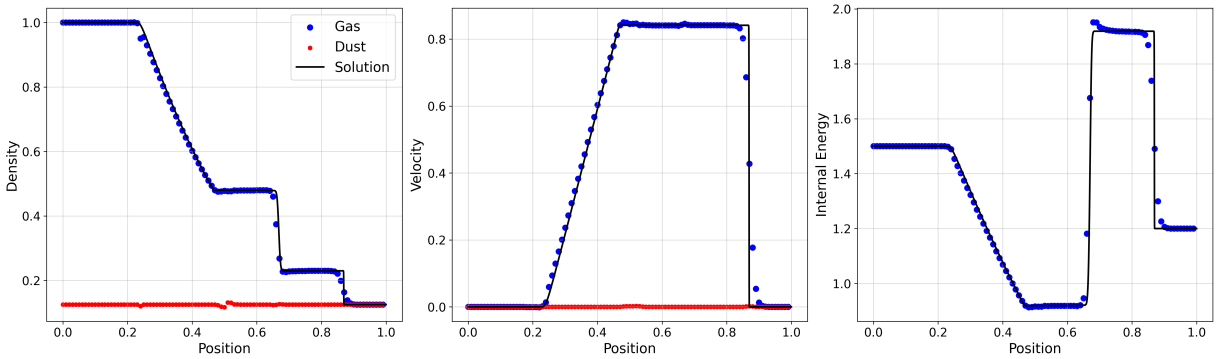

Figure 3.4: DUSTYSHOCK benchmark test for zero drag. The blue and red dots represent the output of the simulation for gas and dust, respectively, while the continuous line shows the solution for the gas evolution. As expected, while the gas undergoes the typical Sod shock tube evolution, the dust remains unperturbed.

The dust, instead, is initialized as a constant background density at rest

$$\rho_L^d = 0.125 \quad \rho_R^d = 0.125,$$
$$\vec{v}_L^d = 0 \quad \vec{v}_R^d = 0.$$

The drag coefficient is set to $K = 0$. The purpose of this test is to verify that the One-Fluid model implementation is able to evolve the gas shock without perturbing the dust background uniform density, which should remain unperturbed given the zero drag it feels from the gas.
The simulation is initialized with a total of $N = 1375$ particles in a domain $x \in (-0.5, 0.5)$. Boundary particles are placed outside each end of the domain to keep the particles confined. The internal energy is initialized and evolved following an ideal gas law, with constant $\gamma = 5/3$.
The results are shown in Fig. 3.4. As expected, the gas undergoes the standard shock as if the dust were not present, and the dust, conversely, stays unperturbed. This excellent performance in keeping the dust still is to be attributed to the Non-Conservative artificial viscosity scheme, which is here fully active in the shocked regions of the fluid.

Strong coupling

We also explore the opposite scenario, a Sod shock tube test, where dust is strongly coupled to the gas. While the gas initial left and right states are identical to the ones used for the weak coupling case, the dust is now initialized using the same density distribution as the gas, as opposed to the constant background used previously

$$\rho_L^d = 1 \quad \rho_R^d = 0.125,$$
$$\vec{v}_L^d = 0 \quad \vec{v}_R^d = 0.$$

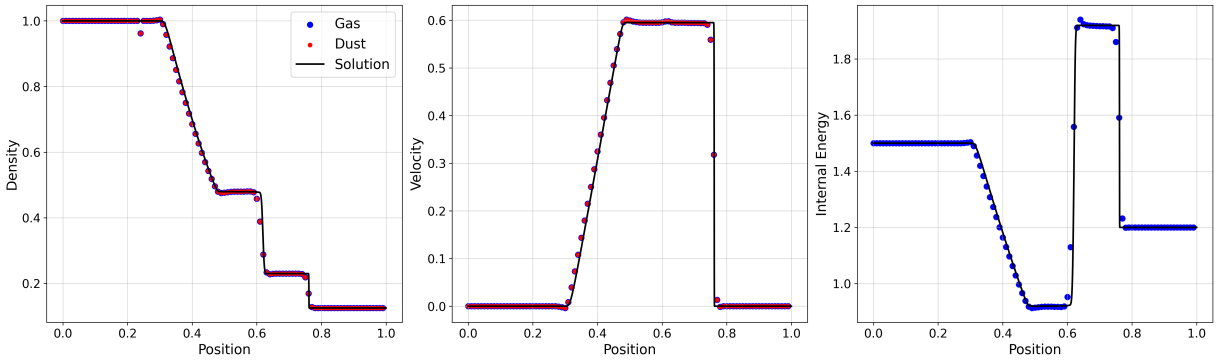

Figure 3.5: DUSTYSHOCK benchmark test for strong drag ($K = 10^6$). The blue and red dots represent the output of the simulation for gas and dust, respectively, while the continuous line shows the solution for the gas evolution. The evolution of the dust now closely follows the gas, and both evolve at a lower speed compared to a gas-only simulation, with the modified sound speed shown in Eq. 3.86.

The drag coefficient is set to $K = 10^6$ and a total of $N = 1700$ particles are employed. The expected behavior for this shock is to evolve as a standard Sod shock tube, but with a modified sound speed of

$$\tilde{c_s} = \frac{c_s}{\sqrt{1-\epsilon}}. \tag{3.86}$$

As can be shown in Fig. 3.5, the result from the simulation follows closely the expected solution, and dust and gas remain tightly coupled.

Two-dimensional shocks

In order to test the One-Fluid model implementation in two dimensions, we employed a strategy discussed in Toro (2009) to turn one-dimensional shocks into two-dimensional radially symmetric shocks. In particular, Toro (2009) shows that the radial evolution of a shock in multiple dimensions can be replicated in one-dimensional simulations by adding a geometrical source term to the Euler equations

$$\partial_t \vec{U} + \partial_r \vec{F}(\vec{U}) = \vec{S}(\vec{U}), \tag{3.87}$$

where $\vec{U}$ is the vector of conserved variables, $\vec{F}(\vec{U})$ are the corresponding fluxes and the geometrical source term $\vec{S}(\vec{U})$ is defined as

$$\vec{S}(\vec{U}) = -\frac{\alpha}{r}\begin{bmatrix} \rho v \\ \rho v^2 \\ v(E+p) \end{bmatrix}. \tag{3.88}$$

By varying the α parameter we obtain a cylindrical ($\alpha = 1$) or a spherical ($\alpha = 2$) symmetry. With this procedure, we can obtain the expected result for two-dimensional

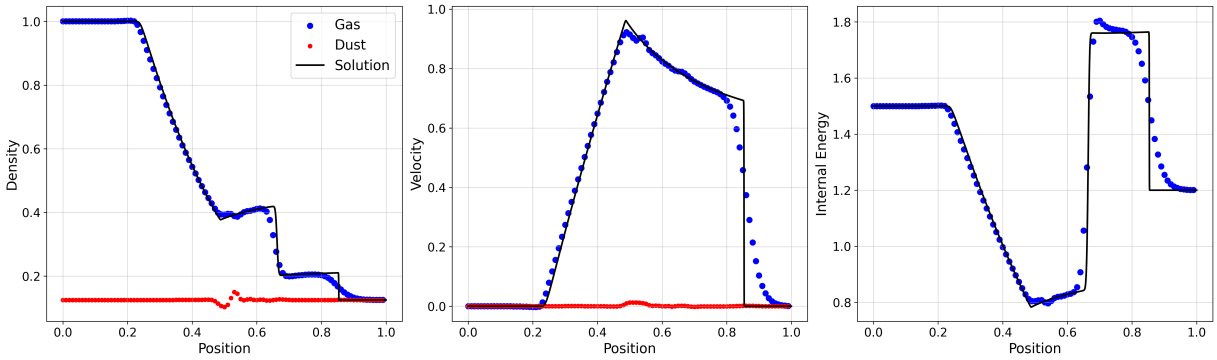

Figure 3.6: Same as in Fig. 3.4, but now in two dimensions. The solution was obtained with a one-dimensional Exact Riemann Solver with an additional geometrical source term as described in Eq. 3.87. The dust now shows some additional wiggles compared to the one-dimensional case, which can be attributed to the lower effective resolution of this higher-dimensional test.

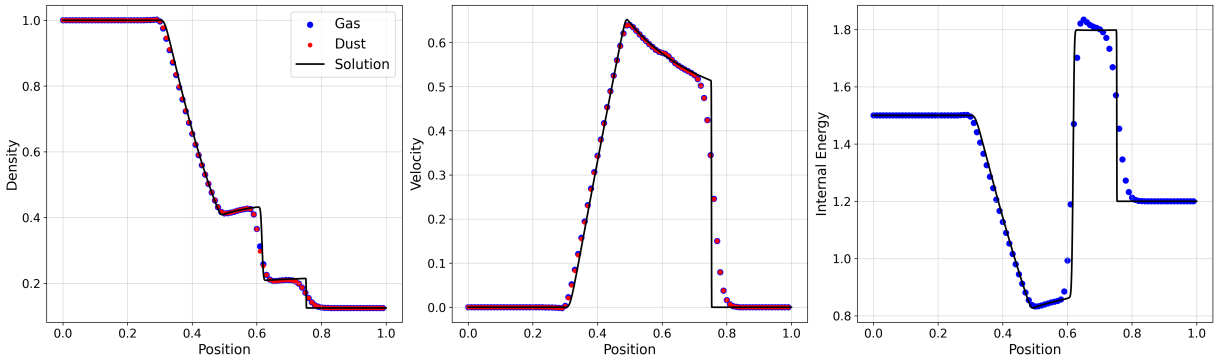

Figure 3.7: Same as in Fig. 3.5, but now in two dimensions. The solution was obtained with a one-dimensional Exact Riemann Solver with an additional geometrical source term as described in Eq. 3.87.

simulations using a one-dimensional Riemann solver.
The initial conditions for the shocks in two dimensions are essentially the same as in the one-dimensional shock. The particles are arranged in a hexagonal closed pack lattice in a periodic domain $x \in (-1, 1)$, $y \in (-1, 1)$. Particles inside a sphere centered around 0 and of radius $R = 0.5$ have the "left" state initial conditions, while the particles outside have the "right" state initial conditions. The results for the two-dimensional shocks for both the weak and strong drag cases are shown in Fig. 3.6 and Fig. 3.7. The performances are slightly worse compared to the one-dimensional case, which can be explained by a lower effective resolution of the multidimensional cases, but the shocks all evolve as expected.

3.3.4 Dusty Sedov-Taylor blast wave

The Sedov-Taylor blast wave is a strong, radially symmetric shock wave, first introduced by Sedov (1946, 1959). It is widely used in astrophysics to model the evolution of supernova blast waves (Steinwandel et al., 2020) and to benchmark the shock-capturing capabilities and limitations of hydrodynamical simulation codes (Springel & Hernquist, 2002; Rosswog & Price, 2007; Groth et al., 2023).
For gas-only simulations, a self-similar solution exists and can be used to verify the performance of numerical schemes. Although no such solution is known for dust–gas mixtures, in the limit of strong drag and high dust-to-gas ratios, the evolution is expected to follow the gas-only solution but at a modified sound speed, as in the *DUSTYSHOCK* tests.
In the context of dust-gas mixtures, Laibe & Price (2012a) simulated a Sedov-Taylor blast wave using a two-fluid model for a moderate drag coefficient of $K = 1$, since a high drag simulation would have been too computationally expensive, and a zero-drag simulation would have been trivial in the two-fluid model. Here we will simulate the Sedov-Taylor blast wave for a mixture of dust and gas in two cases, mirroring the ones used in the DUSTSHOCK tests: (i) a zero-drag simulation with dust-to-gas ratio of 0.01 and (ii) a high-drag ($K = 10^6$) simulation with dust-to-gas ratio of 1.
We initialize $N = 64^3$ particles in a regular grid $x, y, z \in (-0.5, 0.5)$ with periodic boundary conditions, uniform gas density $\rho_g = 1$ and vanishingly low pressure $P = 10^{-6}$, except for the particles closest to the center which are initialized with a large internal energy $U = 10$. Consequently, a shock with Mach number $M \sim 2 \cdot 10^4$ propagates through the domain. The simulation is evolved until $t_\mathrm{f} = 0.02$
In Fig. 3.8 we show the radial density profiles obtained from the simulations for both regimes, zero and high drag, together with the self-similar solution. For $K = 0$ (left), the gas undergoes the Sedov-Taylor blast wave and approaches the expected self-similar solution, while the dust remains largely at rest. For high drag (right), the dust and gas stay tightly coupled, and the mixture can be described by the self-similar solution evaluated at an earlier time, consistent with the modified sound speed given in 3.86.

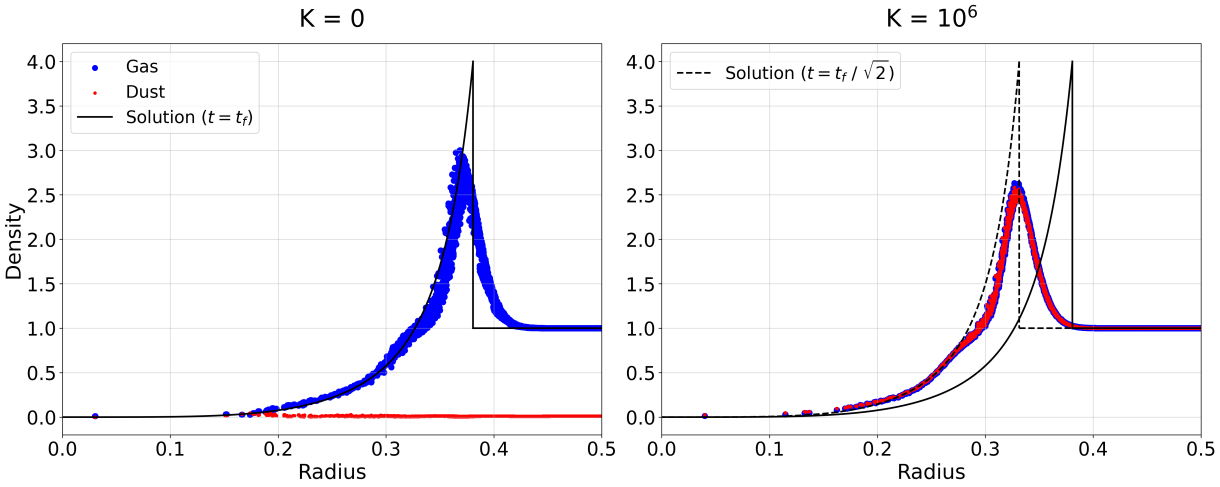

Figure 3.8: Dusty Sedov-Taylor blast wave for a dust and gas mixture in two regimes, $K = 0$ (left) and $K = 10^6$ (right). The radial density profiles are shown as blue dots for gas and red dots for dust, and the self-similar solution at $t = t_\mathrm{f}$ is plotted as a continuous black line. The simulation correctly keeps dust at rest while the gas evolves following the self-similar solution, for uncoupled dust. For strongly coupled gas and dust, the mixture follows the self-similar solution evaluated at a previous time, consistent with the modified sound speed in Eq. 3.86, shown as a dashed black line.

3.3.5 Dusty Cold Keplerian Disk

The Cold Keplerian Disk is a standard hydrodynamical test used to study the stability of numerical frameworks (Cartwright et al., 2009; Hopkins, 2015; Groth et al., 2023). SPH has already been shown to struggle in keeping a Cold Keplerian Disk stable for various orbits; thus, we are not interested in studying its stability. Instead, we will focus on how dust clumping behaves under various dust-to-gas ratios and drag coefficients and how the dust pressure model introduced in Section 2.7 can help keep the code stable.
The initial condition for the Cold Keplerian Disk is taken, for instance, from Groth et al. (2023). The gas surface density is given by

$$\Sigma_g = 0.01 + \begin{cases} \Big(R/0.5\Big)^3 & \text{if } R \leq 0.5, \\ 1 & \text{if } 0.5 \leq R \leq 2, \\ \Big(1 + (R-2)/0.1\Big)^{-3} & \text{if } R > 2. \end{cases} \tag{3.89}$$

Particles are initially distributed in a hexagonal close-packed grid in a domain $x \in (-2, 2)$, $y \in (-2, 2)$, and then removed from $R < 0.5$ and $R > 2$ until the desired density profile is reached. The effective resolution of the initial grid is 256^2. The velocity is set to strictly Keplerian, and the pressure is set at a vanishingly low value of $P = 10^{-6}$. In adding the dust we simply increase the density uniformly by a factor $(1 + \delta)$, where δ is the dust-to-

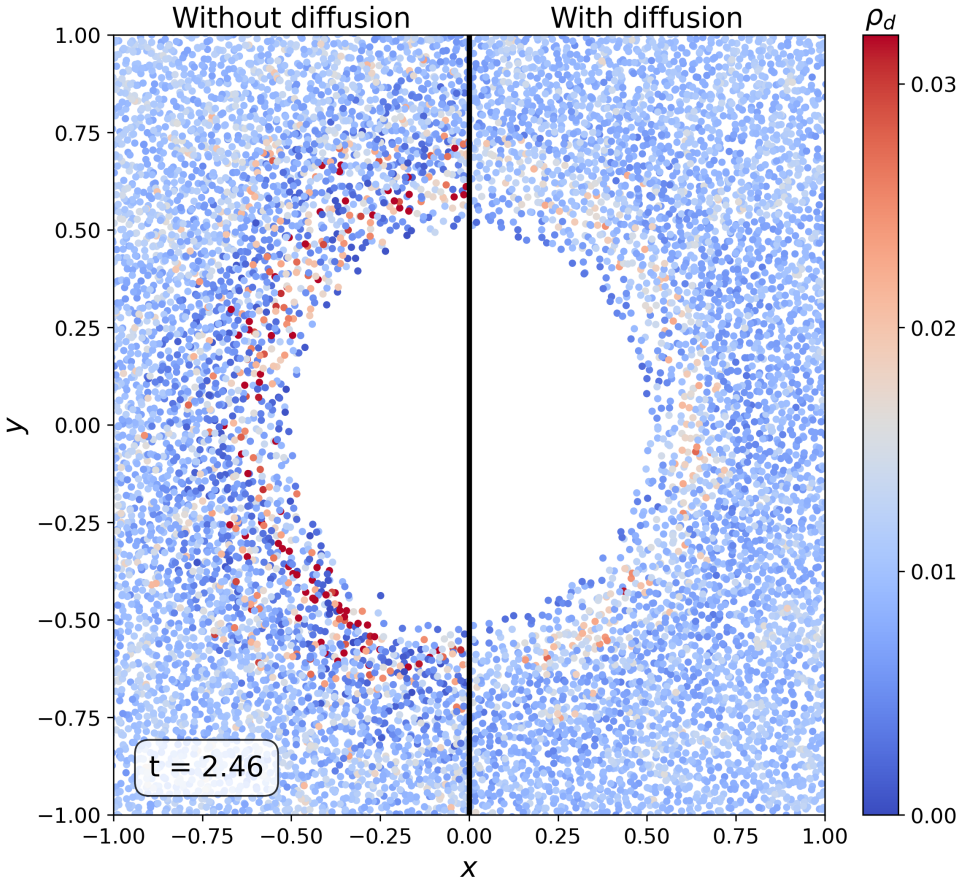

Figure 3.9: Zoom-in in the central region of the Cold Keplerian Disk simulation with weak drag ($K = 0.01$) and low dust-to-gas ratio ($\delta = 0.01$). The left half shows the last snapshot of a run without diffusion. The lack of a restoring force causes strong clumping of dust density, which leads to a code crash very early into the simulation. The addition of diffusion, as shown in the right half of the figure, dissipates these clumps and makes the code more stable.

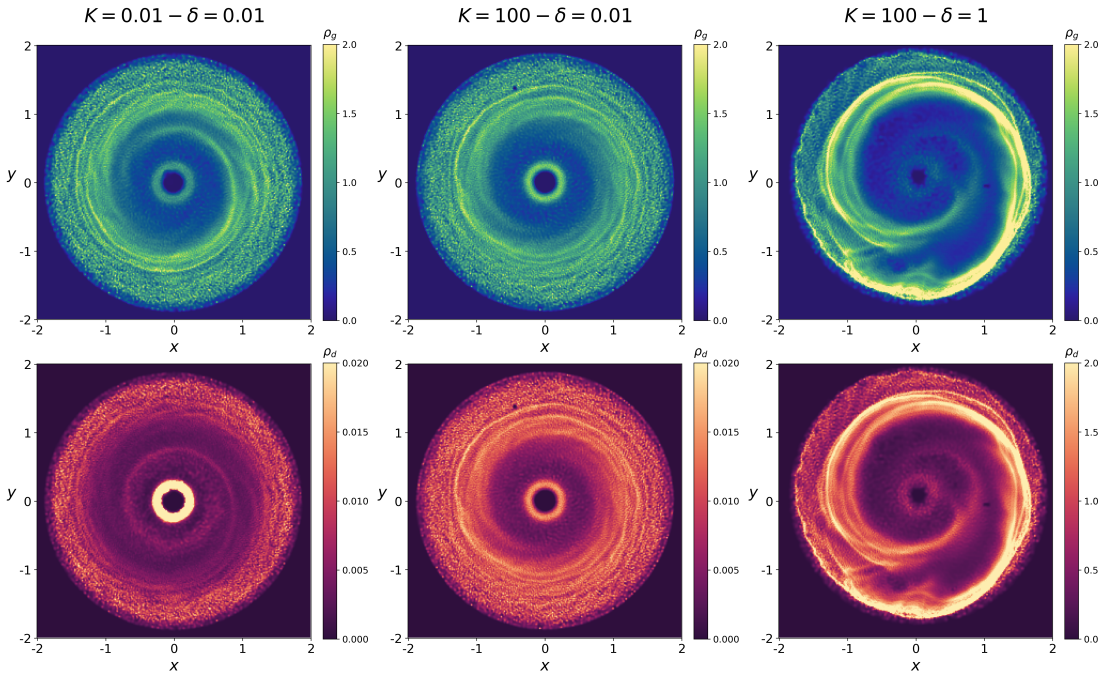

Figure 3.10: Dusty Cold Keplerian Disk simulations for three different sets of parameters. Gas density is shown in the top panels while dust density is shown in the bottom panels.

gas ratio. The dust fraction is set accordingly as a uniform $\epsilon = \delta/(1+\delta)$. Dust is also initialized at a Keplerian velocity, hence, the initial difference velocity vector is $\vec{\Delta v} = 0$. Such initial conditions should remain at equilibrium indefinitely, but in SPH the artificial viscosity eventually causes the disk to break up after roughly 5-10 orbits (Hopkins, 2015). We conduct a series of simulations varying the initial dust-to-gas ratio and the drag coefficient K.

In Fig. 3.9 we show two simulations, both with $\delta = 0.01$ and $K = 0.01$, the only difference being the use of dust pressure as a diffusive term, as described in Section 2.8. The left half of the figure, corresponding to a simulation without any diffusion, shows severe clumping in dust density, which eventually causes a crash in the code. The snapshot shown here is the last available from the simulation. On the right half is shown a simulation where diffusion is switched on, at the same snapshot. As shown in the figure, the dust density field is less clumped, resulting in smoother gradients and a more stable simulation.

Fig. 3.10 presents the results of three distinct simulation configurations at time $t = 120$, corresponding to roughly 20 orbits at $R = 1$. As anticipated, the disk becomes unstable over time in all cases due to the artificial viscosity inherent in SPH. The left panels depict a simulation with weak drag ($K = 0.01$) and a low dust-to-gas ratio ($\delta = 0.01$). In this scenario, the dust gradually drifts inward as it loses angular momentum, but its influence on the gas dynamics is negligible due to the small dust-to-gas ratio. The middle panels illustrate a simulation with strong drag ($K = 100$) and the same low dust-to-gas ratio ($\delta = 0.01$). Here, the dust closely follows the motion of the gas, as expected under

strong coupling conditions. Finally, the right panels show a simulation with strong drag ($K = 100$) and a high dust-to-gas ratio ($\delta = 1$). In this case, the disk undergoes significant disruption, as the pressure gradient force of the gas is insufficient to counteract the influences from the dust dynamics. The dust concentrates in spiral structures that actively drag the gas along, resulting in a highly coupled system where the gas motion is strongly influenced by the dust.
While the presence of dust diffusion helped stabilize the code under various conditions, especially for weak drag regimes, the code could not remain stable for the most extreme conditions of weak and intermediate drag ($K = 0.01$ and $K = 1$) for a dust-to-gas ratio of unity ($\delta = 1$). Although our code aims at being as general-purpose as possible, such extreme conditions are rarely encountered in common astrophysical applications and would require specialized techniques to be simulated.

3.3.6 Dusty Protoplanetary Disk

The Cold Keplerian Disk is a useful and challenging hydrodynamical test to investigate the stability of our implementation. However, a more standard and physically motivated setup is a model of an actual protoplanetary disk. Our configuration follows previous SPH studies of protoplanetary disks (Laibe & Price, 2012b; Vericel et al., 2021), but restricted here to two dimensions.
The surface density distribution follows a power law with a smooth inner edge

$$\Sigma(R) = \Sigma_0 R^{-p} \left(1 - \sqrt{\frac{R_{\rm in}}{R}}\right), \tag{3.90}$$

with $p = 3/2$ and $R_{\rm in} = 10$ au, and Σ_0 chosen so that the total mass of the disk is $M_{\rm disk} = 0.01 M_\odot$. To generate a particle distribution with this density profile, we first use inverse sampling of the cumulative distribution function to generate values for the radii, and random values for the ϕ coordinate. To reduce the statistical noise of the resulting point distribution, we apply the Centroidal Voronoi Tessellation (CVT) technique described in Du et al. (1999). The initial particle set is used to construct the corresponding Voronoi tessellation with the `qvoronoi` command from the `Qhull` library (Barber et al., 1996). We then compute the mass centroid of each Voronoi cell as

$$\vec{x}_{j,i+1} = \frac{\int_{V_{j,i}} \vec{x}\Sigma(\vec{x})dV}{\int_{V_{j,i}} \Sigma(\vec{x})dV}, \tag{3.91}$$

where Σ is the target density profile in Eq. 3.90, the j index spans over the points in the distribution and the i denotes the iteration step. The new point distribution $\vec{x}_{j,i+1}$ is then used to draw a new Voronoi tessellation, and the procedure is iterated until convergence. The resulting point distribution is much smoother than the initial one obtained from inverse sampling.
The CVT was used to generate the positions of $N = 10^6$ points between $R_{\rm in} = 10$ au and

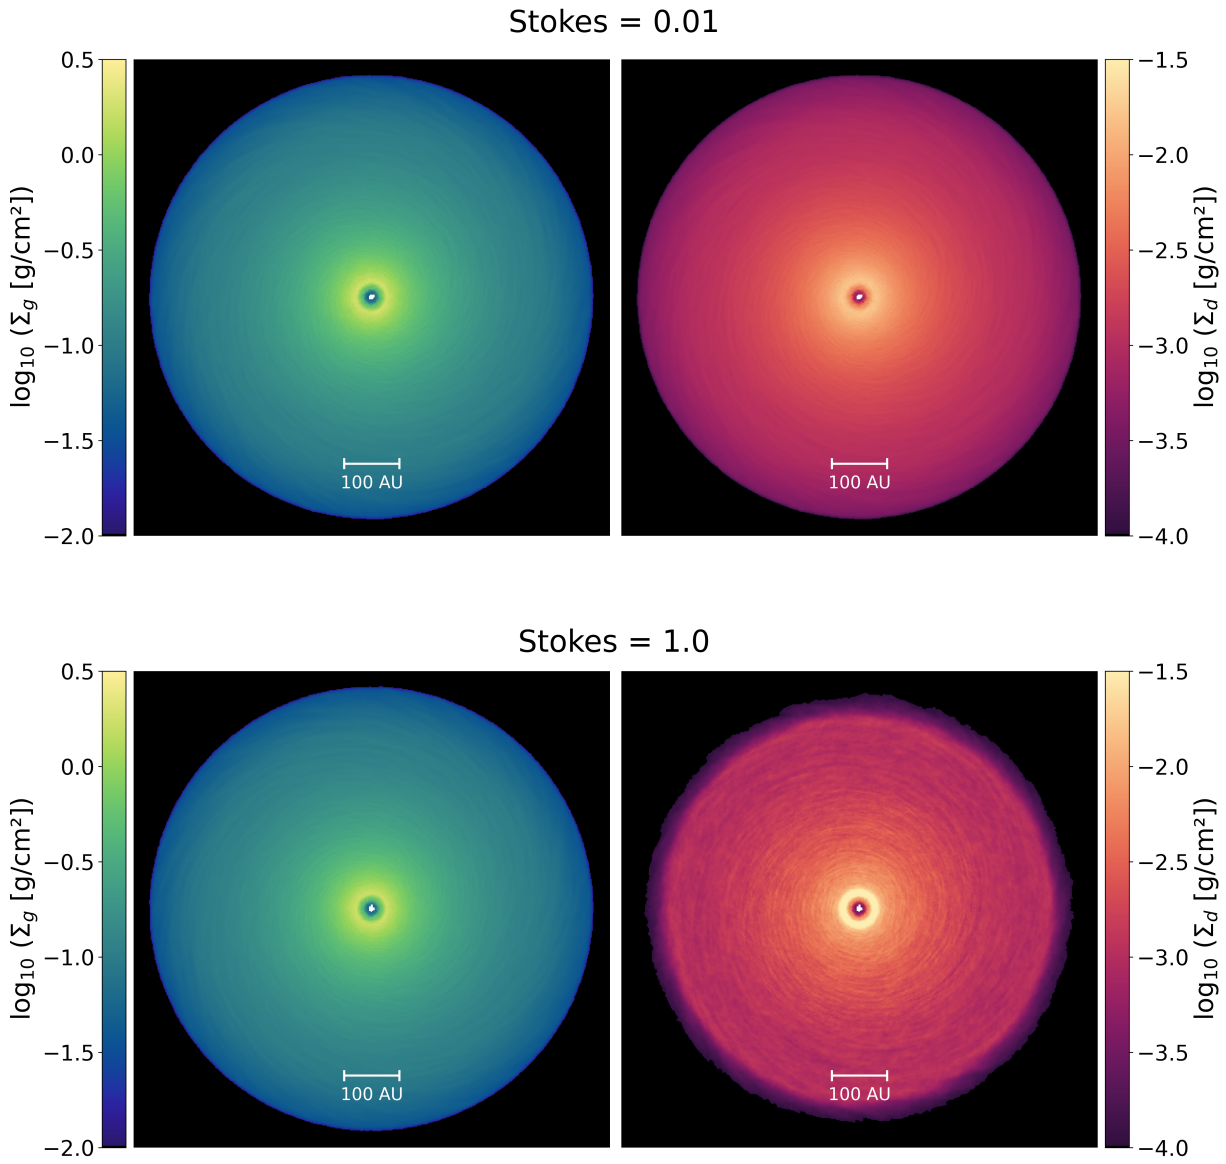

Figure 3.11: Gas (left panels) and dust (right panels) density for a dusty protoplanetary disk for two values of Stokes number at $t = 20\ kyr$. For tightly coupled grains ($St = 0.01$, top row), the gas and dust density remains at equilibrium, while for weakly coupled grains ($St = 1$, bottom row), the dust undergoes both clumping and radial drift toward the pressure maximum.

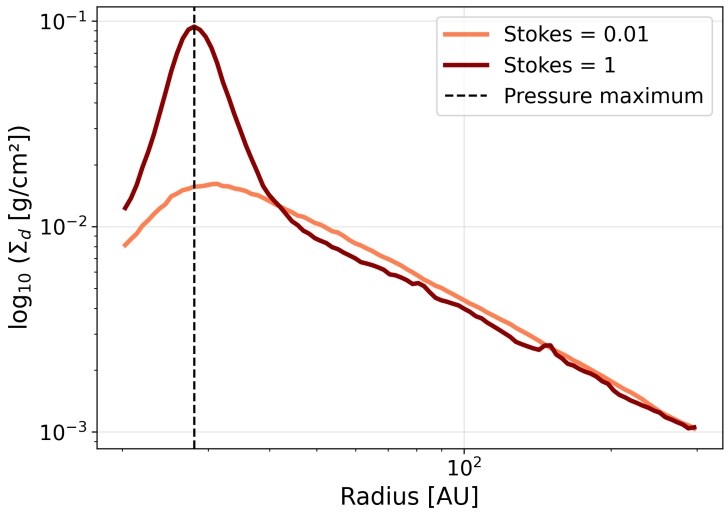

Figure 3.12: Dust radial density profile for both runs with $St = 0.01$ and $St = 1$. The coupled grains keep their original density profile, driven by the strong drag force, which forces them to follow the dynamics of the gas. The weakly coupled grains, on the other hand, undergo a much stronger radial drift toward the position of the pressure maximum.

$R_{\rm out} = 400$ au. The dust-to-gas ratio is set to 0.01, and both gas and dust follow the same initial density profile. The disk is locally isothermal, with the sound speed following the power-law

$$c_s = c_{s,0} R^{-q/2}, \tag{3.92}$$

with the index $q = 3/4$ and the sound speed at radius $R = 1$ au equal to the aspect ratio: $c_{s,0} = h = 0.0281$, as in Vericel et al. (2021).
The azimuthal velocity is set to exactly Keplerian for dust, and slightly sub-Keplerian for the gas

$$v_{g,\phi} = v_K \sqrt{1 - h^2(p+q)}, \tag{3.93}$$

where v_K is the Keplerian azimuthal velocity.
The dust drag coefficient is set through its Stokes number, and two values are used: $St = 0.01$ and $St = 1$, to study the different evolution of tightly and weakly coupled dust grains.
These initial conditions are evolved until $t_{\rm max} = 20\ kyr$ under the influence of the gravitational potential of the central star of mass $M_* = 1\ M_\odot$. During the evolution, particles that cross the inner radius $R_{\rm in}$ are removed from the simulation.
The gas and dust density maps for both St = 0.01 and St = 1 at $t = t_{\rm max}$ are shown in Fig. 3.11. For tightly coupled grains, both the gas and dust density profiles remain stable for long timescales. The pressure support of the gas and the drag force acting on the dust make this test much more stable than the Cold Keplerian Disk counterpart. For $St = 1$, the low dust content is not enough to modify the gas dynamics. However, dust undergoes more significant clumping and radial drift. In Fig. 3.12 we compare the radial dust density

profiles for the two runs, and it is evident that the $St = 1$ run favors radial drift toward the pressure maximum created by the inner radius tapering in Eq. 3.90.

3.3.7 Dusty Kelvin-Helmholtz Instability

To explore the capability of the One Fluid model in simulating the dynamics of dust in turbulent motion, we simulate a Kelvin-Helmholtz instability for various choices of dust-to-gas ratios and drag coefficients. As before, we add the dust to a gas-particle setup by uniformly multiplying the density by $(1+\delta)$. We use the same initial conditions as, for instance, McNally et al. (2012), Hopkins (2015) or Groth et al. (2023): a periodic box in two dimensions is filled with two fluids of density $\rho_1 = 1$ and $\rho_2 = 2$ with opposite velocities $\vec{v}_1 = 0.5\hat{x}$ and $\vec{v}_2 = -0.5\hat{x}$ and a small perturbation in the $\hat{y}$ direction to trigger the instability. We expect the instability to grow with a timescale $\tau_{KH} = \frac{\lambda}{v_{x,1}-v_{x,2}} \frac{\rho_1+\rho_2}{\sqrt{\rho_1\rho_2}}$. In Fig. 3.13 we show both the gas and dust densities at a time $t \simeq 1.2\tau_{KH}$ for the same dust-to-gas ratios and drag coefficients that were shown in Fig. 3.10 for the Cold Keplerian Disk.

In the left panels, corresponding to low dust-to-gas ratio and weak drag, gas is not affected

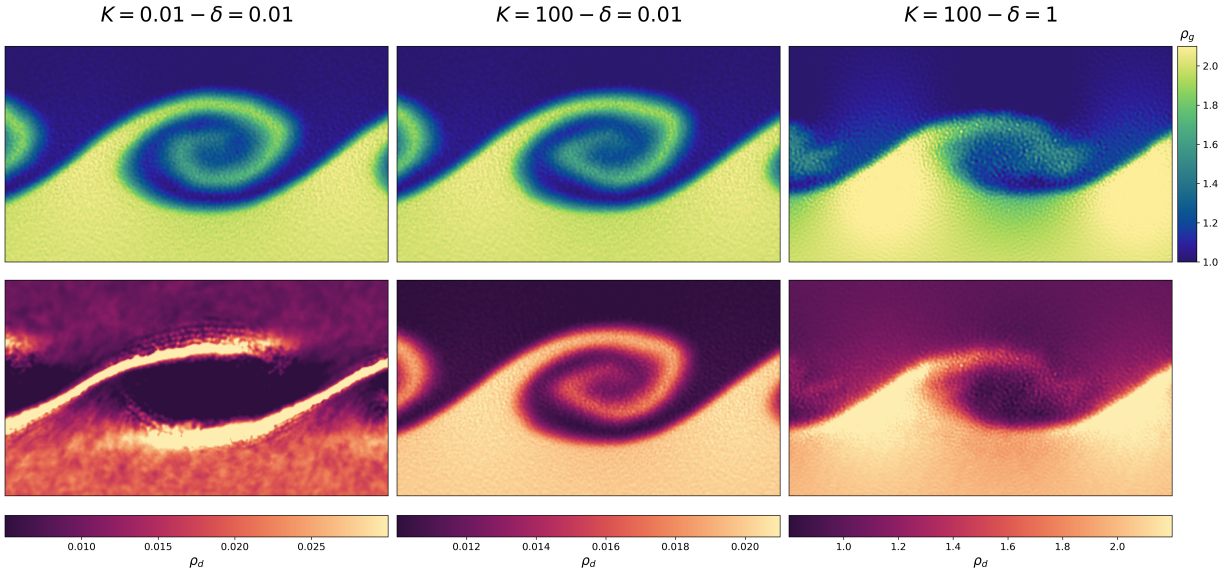

Figure 3.13: Dusty Kelvin-Helmholtz instability simulations for three different sets of parameters. Gas density is shown in the top panels while dust density is shown in the bottom panels.

by the presence of dust, while the dust feels just enough drag to cluster along the edges of the vortex without tracing the gas inside. In the middle panels (low dust-to-gas ratio and high drag), the small amount of dust is stirred so effectively by the gas that it closely follows the evolution of a gas-only simulation. For high drag and high dust-to-gas ratio (right panels), the growth of the instability is slowed and dampened considerably, as gas has

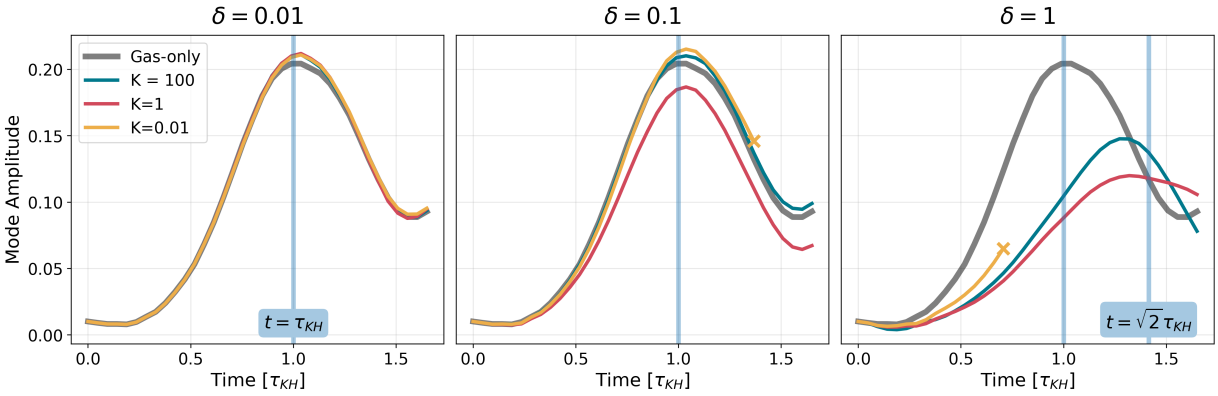

Figure 3.14: Time evolution of the Kelvin-Helmholtz instability mode amplitude for three dust-to-gas ratios, $\delta = (0.01, 0.1, 1)$, and three drag coefficients, $K = (0.01, 1, 100)$. Time is shown in units of the growth timescale τ_{KH}.

to drag more material upward in order to develop its turbulent motion. Both the clustering of dust around the edges of the vortices for low drag and the slowing down of the growth rate for large dust-to-gas ratios were already reported by Hendrix & Keppens (2014), confirming the equivalence of our SPH implementation with their Eulerian approach.
To explore in more detail the growth of the instability for various parameters, we compute the mode amplitude M using discrete convolutions, as described by McNally et al. (2012). The evolution of the mode amplitude is shown in Fig. 3.14 for all dust-to-gas ratios and drag coefficients. In the left panel ($\delta = 0.01$), all curves closely follow the evolution of a gas-only simulation, with the mode amplitude peaking, as expected, at time $t = \tau_{KH}$. Variations in mode amplitude between simulations with different drag coefficients start appearing for $\delta = 0.1$ (middle panel). In particular, the growth for K = 1 is damped, similar to what we observed for the dusty wave at intermediate drag regimes (i.e., where both gas and dust waves were significantly damped). For large dust-to-gas ratio ($\delta = 1$, right panel) the evolution is very different: the growth timescale increases and the amplitude of the growth gets dampened. Considering that the increased dust-to-gas ratio reduces the effective sound speed of the medium, as described in Eq. 3.86, a crude approximation for the new growth timescale is $\tau'_{KH} = \sqrt{2}\tau_{KH}$.
Note that the weak drag ($K = 0.01$) simulations for $\delta = 0.1$ and $\delta = 1$ did not reach the end of the simulation, as for the Cold-Keplerian Disks, because of too large gradients of dust fraction.

3.4 Conclusions

We have described the implementation of the One-Fluid model in `OpenGadget3`, built upon the original formulation and implementation by Laibe & Price (2014b) and Laibe & Price (2014c). Improvements from the original implementations are:

- A correct version of the equation for $\frac{d\vec{\Delta v}}{dt}$, both in the continuum and in its SPH-discretized formulation. The continuum version is presented in two forms: one more compact and one more convenient to discretize, for the energy conservation arguments presented in Section 3. Many other possible forms of this equation can be derived, such as the one presented in Lebreuilly et al. (2019).
- A simple time-dependent extension of artificial viscosity and artificial conductivity, already present in `OpenGadget3` and adapted for a dusty fluid.
- A unified artificial viscosity framework that uses the shock indicator R_a to automatically interpolate between the conservative and non-conservative formulations presented in Laibe & Price (2014c), capitalizing on the benefits from both methods.
- A novel artificial diffusion model for dust, built upon the "pressure-like" term introduced by Klahr & Schreiber (2021), with an adaptive diffusion coefficient.

The implementation was benchmarked against standard dust dynamics tests such as DUSTYBOX, DUSTYWAVE, and DUSTYSHOCK, the last of which was performed in one and two dimensions. Additionally, we further tested the shock-capturing properties of our implementation against a dusty Sedov-Taylor blast wave. All benchmark tests showed excellent agreement with analytic solutions at all drag regimes while conserving mass, momentum, and energy. We additionally demonstrated the robustness of our numerical scheme on two sets of more complex scenarios: Dusty Cold Keplerian Disks, Dusty protoplanetary disks, and Dusty Kelvin-Helmholtz instabilities. Thanks to the dust diffusion model, the code proved to be stable under most parameter choices, the only exception being the weak and intermediate drags ($K = 0.01$ and $K = 1$) for a dust-to-gas ratio of unity. Future improvements will involve implementing the One Fluid model also to the Meshless-Finite-Mass framework in `OpenGadget3` (Groth et al., 2023) as well as extending the implementation to multiple grain sizes.

Acknowledgements

TB acknowledges funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under grant 325594231. GTP, TB, KD, and BE acknowledge support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC-2094 - 390783311. GTP and KD acknowledge support by the COMPLEX project from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program grant agreement ERC-2019-AdG 882679. GTP and TB acknowledge funding from the European Union under the European Union's Horizon Europe Research and Innovation Programme 101124282 (EARLYBIRD). Views and opinions expressed are, however, those of the authors only and do not necessarily reflect those of the European Union or the European Research Council. Neither the

European Union nor the granting authority can be held responsible for them. MH acknowledges support from the DFG program “Closing the Loop - Using Synthetic Observations of Simulated Star-forming Regions to Test Observational Properties” (DFG Project Number: 426714422).

4

General Implicit Runge-Kutta Integrators for Multifluid Gas-Dust Aerodynamic Drag

Published as
Tedeschi-Prades et al., 2025, The Astrophysical Journal, Volume 994, n. 267

Full authors list:
Giovanni Tedeschi-Prades; Til Birnstiel; Klaus Dolag; Barbara Ercolano; Mark Hutchison.

The integration of aerodynamic drag is a fundamental step in simulating dust dynamics in hydrodynamical simulations. We propose a novel integration scheme, designed to be compatible with Strang splitting techniques, which allows for the straightforward integration of external forces and hydrodynamic fluxes in general-purpose hydrodynamic simulation codes. Moreover, this solver leverages an analytical solution to the problem of drag acceleration, ensuring linear complexity even in cases with multiple dust grain sizes, as opposed to the cubic scaling of methods that require a matrix inversion step. This new General Implicit Runge-Kutta integrator (GIRK) is evaluated using standard benchmarks for dust dynamics such as DUSTYBOX, DUSTYWAVE, and DUSTYSHOCK. The results demonstrate not only the accuracy of the method but also the expected scalings in terms of accuracy, convergence to equilibrium, and execution time. GIRK can be easily implemented in hydrodynamical simulations alongside hydrodynamical steps and external forces, and is especially useful in simulations with a large number of dust grain sizes.

4.1 Introduction

Astrophysical dust is a fundamental component of a wide range of objects, from distant galaxies to nearby protoplanetary disks (PPDs). Its thermal emission allows us to probe astrophysical objects in the millimeter-infrared wavelength range, and its presence has an impact on both the thermodynamic and dynamical state of the surrounding gas. Dust grains scatter, absorb, and re-emit infrared radiation, controlling the temperature of the surrounding gas. From a dynamical point of view, dust continuously exchanges momentum with gas via aerodynamic drag, leading to a diverse array of gas-dust interactions that are often non-linear and difficult to predict without detailed numerical simulations. This is particularly true for larger grain sizes that are only partially coupled to the gas or in regions with high dust-to-gas ratios, both common scenarios in PPDs.

Dust dynamics has been the focus of many new developments in modern hydrodynamical simulation codes. In hybrid Eulerian-Lagrangian schemes, gas is evolved on a grid (either static or adaptive), while dust is simulated as a set of "superparticles" that move freely over the mesh (Yang & Johansen, 2016; McKinnon et al., 2018b; Mignone et al., 2019). Other Eurlerian schemes, instead, include dust as an additional fluid alongside gas (Benítez-Llambay et al., 2019; Huang & Bai, 2022). In Lagrangian codes, dust has been treated either as a separate set of particles (Monaghan & Kocharyan, 1995; Laibe & Price, 2012a) or via a 'One-Fluid' approach, where each particle represents a gas–dust mixture rather than a single phase (Laibe & Price, 2014b).

A key aspect of simulating dust dynamics, regardless of the numerical approach, is the integration of the drag term coupling the gas and dust. This drag term can be integrated either explicitly or implicitly; however, explicit integration requires the timestep of the simulation to be smaller than the shortest stopping time from the entire dust population. In many astrophysical systems, the stopping time for the most tightly coupled grains is much shorter than the simulation timestep, rendering explicit integration impractical. For this reason, implicit integration of aerodynamic drag is of paramount importance when simulating the dynamics of gas-dust mixtures.

Several methods have been developed to implicitly integrate the drag for a population of dust sizes, but these methods often require a matrix inversion to be carried out. While matrix inversion algorithms are readily available, these methods can be computationally expensive when a large number of dust sizes are needed, since the cost scales as $\mathcal{O}(N_d^3)$, where N_d is the number of dust species. Krapp & Benítez-Llambay (2020) proposed a closed-form solution describing how velocity evolves under the influence of drag, assuming the momentum transfer between dust species is negligible. This solution only scales linearly with the number of dust grain sizes ($\mathcal{O}(N_d)$). Krapp et al. (2024a) combined this solution with the Diagonally Implicit Runge-Kutta (DIRK) integrator (Ascher et al., 1997; Pareschi & Russo, 2005) to develop the first integration scheme that is second-order accurate with first-order complexity (called Multifluid DIRK, or MDIRK). While MDIRK offers significant computational advantages, its specialized Implicit-Explicit (IMEX) scheme for handling hydrodynamical fluxes and external forces limits its implementation to general-

purpose hydrodynamic codes.
Our work builds upon the analytical solution presented in Krapp & Benítez-Llambay (2020) and the MDIRK integrator. In particular, we present a novel integration scheme, designed to be compatible with Strang splitting techniques, which allows for straightforward integration of external forces and hydrodynamical fluxes in general-purpose hydrodynamical codes. We introduce two versions of this integrator: one tailored to Strang splitting with a single hydrodynamic step, and another for the more commonly used two-step approach in second-order hydrodynamical simulations. This flexibility makes our scheme broadly applicable and easy to implement in a wide range of simulation frameworks.
In Section 4.2, we introduce the general formulation of the problem and briefly review the analytical solution derived by Krapp & Benítez-Llambay (2020). We then develop a systematic procedure for translating desired properties, namely second-order accuracy and asymptotic stability, into a solvable set of algebraic conditions. We demonstrate that our new integration scheme satisfies all of these conditions simultaneously for a particular choice of parameters. In Section 4.3, we benchmark the scheme against a suite of test problems, confirming that it achieves all desired properties. Finally, Section 4.4 discusses the performance of the integrator and outlines potential applications.

4.2 Methods

In this paper, we consider a system consisting of gas and N_d dust species that exchange momentum via aerodynamic drag. We restrict our discussion to a one-dimensional system; however, extensions to two and three dimensions are straightforward, since aerodynamic drag does not couple motions in perpendicular directions and can be treated independently. We denote gas properties with the subscript 'g' and dust properties with the subscript 'd'. Furthermore, an index $i = 1, ..., N_\mathrm{d}$ is used to label the different dust species.
The continuity equations for the gas and the N_d dust species read

$$\begin{cases} \partial_t \rho_\mathrm{g} + \partial_x(\rho_\mathrm{g} v_\mathrm{g}) = 0, \\ \partial_t \rho_{\mathrm{d},i} + \partial_x(\rho_{\mathrm{d},i} v_{\mathrm{d},i}) = 0, \end{cases} \tag{4.1}$$

where ρ and v denote the density and velocity, respectively. The momentum conservation equations are

$$\begin{cases} \partial_t(\rho_\mathrm{g} v_\mathrm{g}) + \partial_x(\rho_\mathrm{g} v_\mathrm{g}^2) = -\partial_x P + G_\mathrm{g} + \sum_{i=1}^{N_\mathrm{d}} \rho_\mathrm{g} \alpha_{\mathrm{g},i}(v_{\mathrm{d},i} - v_\mathrm{g}), \\ \partial_t(\rho_{\mathrm{d},i} v_{\mathrm{d},i}) + \partial_x(\rho_{\mathrm{d},i} v_{\mathrm{d},i}^2) = G_{\mathrm{d},i} - \rho_{\mathrm{d},i} \alpha_{\mathrm{d},i}(v_{\mathrm{d},i} - v_\mathrm{g}), \end{cases} \tag{4.2}$$

where P represents the gas internal pressure, G any external force, and the remaining terms on the right-hand sides account for aerodynamic drag. Here, the momentum exchange rates $\alpha_{\mathrm{d},i}$ are simply the inverses of the stopping times $t_{s,i}$. Enforcing momentum conservation, the gas momentum exchange rate $\alpha_{\mathrm{g},i}$ is related to the dust term via

$$\alpha_{\mathrm{g},i} = \epsilon_i \alpha_{\mathrm{d},i}, \tag{4.3}$$

where $\epsilon_i = \rho_g/\rho_{d,i}$ is the dust fraction of the i-th dust specie. In the following, we will only refer to dust momentum exchange rates and drop the subscript 'd'.
In a more compact notation, the momentum conservation equations can be rewritten as

$$\partial_t \mathbf{u} + \partial_x \mathcal{F} = \mathbf{M}\mathbf{u} + \mathbf{G}, \tag{4.4}$$

where:

$$\mathbf{u} = (\rho_g v_g,\ \rho_{d,1} v_{d,1},\ \ldots,\ \rho_{d,N_d} v_{d,N_d}), \tag{4.5}$$

$$\mathcal{F} = (\rho_g v_g^2 + P,\ \rho_{d,1} v_{d,1}^2,\ \ldots,\ \rho_{d,N_d} v_{d,N_d}^2), \tag{4.6}$$

$$\mathbf{G} = (G_g,\ G_{d,1},\ \ldots,\ G_{d,N_d}), \tag{4.7}$$

$$\mathbf{M} = \begin{pmatrix} -\sum_{k=1}^{N_d} \epsilon_k \alpha_k & \alpha_1 & \alpha_2 & \cdots & \alpha_{N_d} \\ \epsilon_1 \alpha_1 & -\alpha_1 & 0 & \cdots & 0 \\ \epsilon_2 \alpha_2 & 0 & -\alpha_2 & \ddots & \vdots \\ \vdots & \vdots & \ddots & \ddots & 0 \\ \epsilon_{N_d} \alpha_{N_d} & 0 & \cdots & 0 & -\alpha_{N_d} \end{pmatrix}. \tag{4.8}$$

In what follows, we will first focus on numerical solutions to the homogeneous system

$$\partial_t \mathbf{u} = \mathbf{M}\mathbf{u}, \tag{4.9}$$

and later discuss how to incorporate the remaining terms from Eq. 4.4.

4.2.1 Explicit Integration

The simplest approach is to integrate the drag term $\mathbf{M}\mathbf{u}$ explicitly. For example, using a first-order Euler method yields

$$\mathbf{u}^{n+1} = \mathbf{u}^n + \Delta t\, \mathbf{M}\mathbf{u}^n. \tag{4.10}$$

The same explicit integration can also be carried out by higher-order schemes, such as Runge-Kutta integrators. However, explicit methods require the timestep to be restricted to the smallest stopping time in the simulation, often rendering the timestep prohibitively small (Laibe & Price, 2014c).

4.2.2 Implicit Integration

Using an implicit Backward Euler scheme, Eq. 4.9 can be integrated via matrix inversion:

$$\mathbf{u}^{n+1} = (1 - \Delta t\, \mathbf{M})^{-1} \mathbf{u}^n. \tag{4.11}$$

Although this method removes the strict timestep restriction of the explicit scheme, it necessitates inverting a matrix, a process with a computational complexity of $\mathcal{O}(N_d^3)$.
For applications requiring a large number of dust bins, it is advantageous to avoid matrix inversion by employing more direct analytical solutions.

4.2.3 Analytical Solution by Krapp & Benítez-Llambay (2020)

To bypass matrix inversion, Krapp & Benítez-Llambay (2020) presented an analytical solution for the drag integration problem that reduces the complexity from third to first order, assuming that dust species exchange momentum only with the gas (and not among themselves). While their solution is formulated in terms of primitive variables, Krapp et al. (2024a) extended it to conserved variables and to the generalized problem of solving

$$(\mathbf{1} - \gamma\,\Delta t\,\mathbf{M})\mathbf{k} = \mathbf{M}\mathbf{q}, \tag{4.12}$$

for arbitrary $\mathbf{k}$ and $\mathbf{q}$ vectors, each having $N_\mathrm{d}+1$ elements (one for each fluid), and with γ being a free parameter. The solution for the gas component is then

$$k_\mathrm{g}^n = \frac{\mathcal{A} - q_\mathrm{g}\,\mathcal{B}}{1+\gamma\,\Delta t\,\mathcal{B}}, \tag{4.13}$$

and for the i-th dust component

$$k_{d,i} = \frac{\alpha_i}{1+\gamma\,\Delta t\,\alpha_i}\Big(\epsilon_i\,q_\mathrm{g} - q_{\mathrm{d},i} + \gamma\,\Delta t\,\epsilon_i\,k_\mathrm{g}\Big), \tag{4.14}$$

with the constants

$$\mathcal{A} = \sum_{k=1}^{N_\mathrm{d}} \frac{\alpha_k q_{\mathrm{d},k}}{1+\gamma\,\Delta t\,\alpha_k}, \quad \mathcal{B} = \sum_{k=1}^{N_\mathrm{d}} \frac{\epsilon_k\,\alpha_k}{1+\gamma\,\Delta t\,\alpha_k}. \tag{4.15}$$

This method solves Eq. 4.9 without inverting a matrix, thereby reducing the computational complexity from $\mathcal{O}(N_\mathrm{d}^3)$ to $\mathcal{O}(N_\mathrm{d})$.

4.2.4 Diagonally Implicit Runge-Kutta (DIRK) Integration from Krapp et al. (2024a)

The analytical solution in Eqs. 4.12 - 4.15 can be incorporated into a second-order integrator scheme, such as the Diagonally Implicit Runge-Kutta (DIRK) integrator, as was done in Krapp et al. (2024a). In this framework, the integration of Eq. 4.9 is

$$\mathbf{u}^{n+1} = \mathbf{u}^n + \Delta t\,(1-\gamma)\,\mathbf{k}_1 + \Delta t\,\gamma\,\mathbf{k}_2, \tag{4.16}$$

where γ is a free parameter and the vectors $\mathbf{k}_1$ and $\mathbf{k}_2$ are determined from

$$\begin{cases} (1-\gamma\,\Delta t\,\mathbf{M})\mathbf{k}_1 = \mathbf{M}\mathbf{u}_n, \\ (1-\gamma\,\Delta t\,\mathbf{M})\mathbf{k}_2 = \mathbf{M}\Big(\mathbf{u}_n + (1-\gamma)\Delta t\,\mathbf{k}_1\Big). \end{cases} \tag{4.17}$$

Eq. 4.17 presents two equations, both with the same structure as Eq. 4.12. The analytical solution can thus be used to find a solution for both $\mathbf{k}_1$ and $\mathbf{k}_2$. This integrator is termed "diagonally implicit" because $\mathbf{k}_1$ can be computed independently of $\mathbf{k}_2$, allowing sequential computation rather than solving a coupled system. Krapp et al. (2024a) coupled this integrator with an IMEX scheme for external forces and hydrodynamical fluxes, and named the combined method MDIRK. The explicit forms of $\mathbf{k}_1$ and $\mathbf{k}_2$ are provided in Krapp et al., 2024a, (Appendix A).

4.2.5 Imposing Desirable Properties

Before introducing our new method, we reframe the discussion into the more general problem of solving Eq. 4.4, which includes external forces and hydrodynamical fluxes. In numerical simulations, particularly when implicit integration is applied to source terms (such as aerodynamic drag), operator splitting is typically used. Popular choices include Strang splitting (a second-order method) or Lie-Trotter splitting (a first-order method where operators are applied for a full step consecutively). Other possible solutions are mixed IMEX methods, which require more specific time-stepping schemes, like the MDIRK integrator (Krapp et al., 2024a).
Ideally, we want the full integration scheme used to solve Eq. 4.4 to preserve the second-order accuracy of the drag integrator. Additionally, we aim to enforce second-order convergence in the presence of external forces, a property that is not guaranteed in general. This was already discussed in Appendix C of Krapp et al. (2024a), where the authors describe how their DIRK integrator, coupled with a Strang splitting scheme, fails to both achieve second-order accuracy and to converge to the correct equilibrium in the presence of external forces.
In the following, we introduce a framework for assessing second-order convergence and accuracy in a given integration scheme and show how any DIRK integrator, when combined with a Strang splitting scheme, fails to satisfy both simultaneously.
We begin by defining the hydrodynamical step operator, $\mathcal{H}_{\Delta t}$, which applies an explicit integration of the hydrodynamical fluxes and external force terms to the state vector $\mathbf{u}$. Since our focus here is on the convergence properties of the scheme in the presence of a constant external force, we restrict $\mathcal{H}_{\Delta t}$ to the simplified form

$$\mathcal{H}_{\Delta t}\mathbf{u} = \mathbf{u} + \mathbf{G_0}\Delta t, \tag{4.18}$$

where $\mathbf{G_0}$ is a constant external force.
Similarly, we define the drag operator $\mathcal{D}_{\Delta t}$, whose action on the state vector $\mathbf{u}$ depends on the chosen implicit integrator. For instance, the DIRK integrator (Krapp et al., 2024a) can be recast as

$$\mathcal{D}_{\Delta t}\mathbf{u} = \Big(1 + \mathbf{A}\Delta t + \gamma(1-\gamma)\mathbf{A}^2\Delta t^2\Big)\mathbf{u}, \tag{4.19}$$

where $\mathbf{A} = (1 - \gamma\,\Delta t\,\mathbf{M})^{-1}\mathbf{M}$.
These two operators can be combined using a Strang splitting scheme, such as

$$\mathbf{u}^{n+1} = \mathcal{D}_{\Delta t/2}\,\mathcal{H}_{\Delta t}\,\mathcal{D}_{\Delta t/2}\,\mathbf{u}^n. \tag{4.20}$$

To facilitate analysis, we decouple the system by diagonalizing the matrix $\mathbf{M}$. Let $\mathbf{P}$ be the transformation matrix that diagonalizes $\mathbf{M}$. Applying $\mathbf{P}^{-1}$ to Eq. 4.20 yields

$$\hat{u}_k^{n+1} = \left(\mathbf{P}^{-1}\,\mathcal{D}_{\Delta t/2}\,\mathcal{H}_{\Delta t}\,\mathcal{D}_{\Delta t/2}\,\mathbf{u}^n\right)_k, \tag{4.21}$$

where $\hat{u}_k^{n+1}$ is the k-th entry of the eigenspace-projected vector $\hat{\mathbf{u}}^{n+1} = \mathbf{P}^{-1}\mathbf{u}^{n+1}$, and corresponds to the eigenvalue λ_k of $\mathbf{M}$. The zeroth component of $\hat{\mathbf{u}}^{n+1}$ corresponds to the

bulk motion of the gas-dust mixture. All other components correspond to the velocity differences between gas and individual dust species, and are the focus of our analysis.
With the system decoupled, we can expand Eq. 4.21 in series and impose convergence and accuracy conditions. Importantly, we will need to perform the expansions both for $\Delta t \to 0$ and $\Delta t \to \infty$, to verify consistency with the differential equation's exact solution and ensure asymptotic stability in the stiff regime, respectively.
In the limit $\Delta t \to 0$, the expansion of the numerical update in Eq. 4.21 will contain terms proportional to $\hat{u}_k^n$ as well as terms proportional to the projected external force $\hat{G}_k$

$$\begin{aligned}\hat{u}_k^{n+1} \simeq \hat{u}_k^n[1 + A_1\Delta t + A_2\Delta t^2 + \mathcal{O}(\Delta t^3)]+ \\ + \hat{G}_k[B_1\Delta t + B_2\Delta t^2 + \mathcal{O}(\Delta t^3)].\end{aligned} \tag{4.22}$$

To impose second-order convergence, we need to compare this against the exact solution of the diagonalized version of Eq. 4.4

$$\frac{d\hat{u}_k}{dt} = \lambda_k \hat{u}_k + \hat{G}_k. \tag{4.23}$$

Expanding it in Taylor series, we obtain the following expected solution for $\Delta t \to 0$

$$\begin{aligned}\hat{u}_k(t+\Delta t) \simeq \hat{u}_k(t)[1 + \lambda_k\Delta t + \frac{1}{2}\lambda_k^2\Delta t^2 + \mathcal{O}(\Delta t^3)]+ \\ + \hat{G}_k[\Delta t + \frac{1}{2}\lambda_k\Delta t^2 + \mathcal{O}(\Delta t^3)].\end{aligned} \tag{4.24}$$

Comparing Eqs. 4.22 and 4.24 we obtain constraints on A_1, A_2, B_1, and B_2 that ensure second-order convergence for $\Delta t \to 0$. Higher-order schemes can be constructed by matching higher-order terms in the expansion.
For the opposite limit, $\Delta t \to \infty$, the process is similar. We first expand the numerical update in Eq. 4.21

$$\begin{aligned}\hat{u}_k^{n+1} \simeq \hat{u}_k^n[C_0 + C_1\Delta t^{-1} + \mathcal{O}(\Delta t^{-2})]+ \\ + \hat{G}_k[D_0 + D_1\Delta t^{-1} + \mathcal{O}(\Delta t^{-2})].\end{aligned} \tag{4.25}$$

In the absence of external forces, the expected behavior for $k \neq 0$ is $\hat{u}_k^{n+1} \to 0$ as $\Delta t \to \infty$, since the velocity difference between gas and dust is damped out. Hence, for second-order convergence of the drag terms, we require $C_0 = 0$ and $C_1 = 0$, so that the leading term in the error is of second order. With external forces, the equilibrium solution is $\hat{u}_k^{n+1} \to \hat{G}_K/|\lambda_k|$ (Krapp et al., 2024a), which constrains the value of D_0. To ensure second-order convergence to this equilibrium, we additionally require $D_1 = 0$.
In practice, it is often not possible to find a single set of parameters that satisfies all constraints in both limits. However, we can switch between two parameter sets depending on whether the timestep is smaller or larger than the maximum stopping time $t_{\rm s}^{\rm max}$. For instance, the DIRK integrator, coupled with the Strang splitting scheme in Eq. 4.20, is

able to satisfy nearly all our desired properties with the following choices of γ

$$\gamma = \begin{cases} 1 \pm \dfrac{1}{\sqrt{2}} & \Delta t < t_{\mathrm{s}}^{\max}, \\ 2 \pm \sqrt{2} & \Delta t > t_{\mathrm{s}}^{\max}. \end{cases} \tag{4.26}$$

The only unsatisfied condition is second-order convergence to the equilibrium solution in the presence of external forces (i.e., $D_1 \neq 0$). To address this, we must extend the approach of Krapp et al. (2024a) to a more general, non-diagonal, implicit Runge–Kutta integrator.

4.2.6 General Implicit Runge-Kutta (GIRK) Integrator

Since the DIRK integrator cannot satisfy all the desired properties simultaneously, we extend our approach to a more General Implicit Runge-Kutta (GIRK) integrator. Starting from the backbone of the Runge-Kutta integration scheme, new integrators can be constructed by constraining the free parameters in the Butcher's Table according, for instance, to the desired integration order or the stability properties or the performances for a specific problem desired. While some general methods have been developed to generate new Runge-Kutta integrators (Maurya et al., 2021; Mizerová & Tvrdá, 2024), we will here derive our own set of parameters that can satisfy all the properties introduced in Section 4.2.5.

The update step of a GIRK is of the form

$$\mathbf{u}^{n+1} = \mathbf{u}^n + \Delta t b \mathbf{k}_1 + \Delta t (1-b) \mathbf{k}_2. \tag{4.27}$$

Here, $\mathbf{k}_1$ and $\mathbf{k}_2$ are obtained by solving the coupled system

$$\begin{cases} (1 - \gamma_1 \Delta t\, \mathbf{M})\, \mathbf{k}_1 = \mathbf{M}\mathbf{u}_n + \Delta t\, \beta_1\, \mathbf{M}\, \mathbf{k}_2, \\ (1 - \gamma_2 \Delta t\, \mathbf{M})\, \mathbf{k}_2 = \mathbf{M}\mathbf{u}_n + \Delta t\, \beta_2\, \mathbf{M}\, \mathbf{k}_1. \end{cases} \tag{4.28}$$

This scheme introduces five parameters: γ_1, γ_2, β_1, β_2, and b. Although the resulting expressions are more involved than in the DIRK case, they still admit an analytical solution. The gas components are given by

$$k_{1\mathrm{g}} = \frac{\beta_1 \Delta t\, C_1 (A_2 - B_2 u_{\mathrm{g}}) - D_2 (A_1 - B_1 u_{\mathrm{g}})}{\beta_1 \beta_2 \Delta t^2 C_1 C_2 - D_1 D_2}, \tag{4.29}$$

$$k_{2\mathrm{g}} = \frac{\beta_2 \Delta t\, C_2 (A_1 - B_1 u_{\mathrm{g}}) - D_1 (A_2 - B_2 u_{\mathrm{g}})}{\beta_1 \beta_2 \Delta t^2 C_1 C_2 - D_1 D_2}, \tag{4.30}$$

and the dust components for the i-th specie are

$$\begin{aligned} k_{1\mathrm{d},i} = \alpha_i\, \Lambda_i \Big[& (u_{\mathrm{g}} \epsilon_i - u_{\mathrm{d},i}) \Big(1 + \alpha_i \Delta t (\gamma_2 - \beta_1) \Big) \\ & + k_{1\mathrm{g}}\, \epsilon_i \Delta t \Big(\gamma_1 + \alpha_i \Delta t\, (\gamma_1 \gamma_2 - \beta_1 \beta_2) \Big) + k_{2\mathrm{g}}\, \beta_1 \epsilon_i \Delta t \Big], \end{aligned} \tag{4.31}$$

$$k_{2\mathrm{d},i} = \alpha_i \Lambda_i \Big[(u_\mathrm{g}\epsilon_i - u_{\mathrm{d},i}) \Big(1 + \alpha_i \Delta t (\gamma_1 - \beta_2) \Big) + k_{2\mathrm{g}}\, \epsilon_i \Delta t \Big(\gamma_2 + \alpha_i \Delta t\, (\gamma_1\gamma_2 - \beta_1\beta_2) \Big) + k_{1\mathrm{g}}\, \beta_2 \epsilon_i \Delta t \Big]. \tag{4.32}$$

All the constants used in these equations are defined in Appendix 4.4. Despite the added complexity, this GIRK integrator remains linear in computational cost, as it avoids matrix inversion and instead relies on the analytical expressions derived in Krapp et al. (2024a). The additional parameters also introduce more degrees of freedom, which can be used to enforce the full set of desirable properties discussed earlier. The drag operator associated with the GIRK operator can be written as

$$\mathcal{D}_{\Delta t}\mathbf{u} = \Big(1 - \Delta t^2\, \beta_1\beta_2\, \mathbf{A}_1\, \mathbf{A}_2 \Big)^{-1} \Big(1 + b\, \Delta t\, \mathbf{A}_1 + (1-b)\, \Delta t\, \mathbf{A}_2 + \Big[\beta_2(1-b) + b\beta_1 - \beta_1\beta_2 \Big] \Delta t^2\, \mathbf{A}_1\, \mathbf{A}_2 \Big), \tag{4.33}$$

where $\mathbf{A}_1 = (1 - \gamma_1 \Delta t\, \mathbf{M})^{-1}\mathbf{M}$ and similarly for $\mathbf{A}_2$.

4.2.7 Strang-Splitting Schemes and Parameter Values

We now propose two Strang splitting schemes to couple GIRK within hydrodynamical simulation codes. The first, simpler scheme is the one already described as

$$\mathbf{u}^{n+1} = \mathcal{D}_{\Delta t/2}\, \mathcal{H}_{\Delta t}\, \mathcal{D}_{\Delta t/2}\, \mathbf{u}^n. \tag{4.34}$$

Having more free parameters, compared to the DIRK integrator, we are able to enforce additional properties. For instance, expanding it for $\Delta t \to 0$, we find that third-order convergence can be achieved with the following parameter choices[1]

$$\begin{cases} b = 1, \\ \beta_1 = \dfrac{1}{2} - \gamma_1, \\ \beta_2 = \dfrac{1 - 3\gamma_1 - 3\gamma_2 + 6\gamma_1\gamma_2}{3 - 6\gamma_1}, \end{cases} \qquad \text{if } \Delta t < t_\mathrm{s}^\mathrm{max}. \tag{4.35}$$

The values of γ_1 and γ_2 are still unconstrained, and can be tuned to better performances by minimizing the errors in benchmark tests. The values we suggest after some testing are $\gamma_1 = 1$ and $\gamma_2 = 0$.

[1]The calculations were carried out using a `Mathematica` notebook that can be downloaded at the following link: https://github.com/Giovanni-Tedeschi/GIRK_MathematicaNotebook

A similar procedure in the opposite limit leads to:

$$\begin{cases} b = 0, \\ \gamma_2 = 2 - \gamma_1, \\ \beta_1 = \dfrac{2 - 2\gamma_1 + \gamma_1^2}{2 - \gamma_1}, \\ \beta_2 = \gamma_1 - 2, \end{cases} \qquad \text{if } \Delta t > t_{\mathrm{s}}^{\max}, \tag{4.36}$$

where the γ_1 is, again, unconstrained. We propose, after some testing, $\gamma_1 = 1.0$.
The second Strang splitting scheme we propose is designed to include two hydrodynamical steps per update, as this is usually the case in most second-order hydrodynamical simulation codes. One possibility to derive such an algorithm is to chain the previous scheme twice

$$\mathbf{u}^{n+1} = \mathcal{D}_{\Delta t/4}\, \mathcal{H}_{\Delta t/2}\, \mathcal{D}_{\Delta t/4}\, \mathcal{D}_{\Delta t/4}\, \mathcal{H}_{\Delta t/2}\, \mathcal{D}_{\Delta t/4}\, \mathbf{u}^n. \tag{4.37}$$

Alternatively, we can optimize this splitting scheme by combining the two central $\mathcal{D}_{\Delta t/4}$ operators

$$\mathbf{u}^{n+1} = \mathcal{D}_{\Delta t/4}\, \mathcal{H}_{\Delta t/2}\, \mathcal{D}_{\Delta t/2}\, \mathcal{H}_{\Delta t/2}\, \mathcal{D}_{\Delta t/4}\, \mathbf{u}^n. \tag{4.38}$$

This new splitting scheme has the same convergence and accuracy scalings for $\Delta t \to 0$ as the previous one, and can be used with the same parameters as in Eq. 4.35. In the stiff limit, however, third-order convergence can be achieved with the following parameters

$$\begin{cases} b = 1, \\ \gamma_2 = 3 - \gamma_1, \\ \beta_1 = -\gamma_1 - 1, \\ \beta_2 = \dfrac{\gamma_1^2 - 3\gamma_1 + 4}{1 + \gamma_1}, \end{cases} \qquad \text{if } \Delta t > t_{\mathrm{s}}^{\max}, \tag{4.39}$$

with $\gamma_1 = 1$. Table 4.1 summarizes the recommended parameter values for both splitting schemes and both timestep regimes.

	$\mathcal{D}_{\Delta t/2}\, \mathcal{H}_{\Delta t}\, \mathcal{D}_{\Delta t/2}$		$\mathcal{D}_{\Delta t/4}\, \mathcal{H}_{\Delta t/2}\, \mathcal{D}_{\Delta t/2}\, \mathcal{H}_{\Delta t/2}\, \mathcal{D}_{\Delta t/4}$	
	$\Delta t < t_s^{\max}$	$\Delta t > t_s^{\max}$	$\Delta t < t_s^{\max}$	$\Delta t > t_s^{\max}$
γ_1	1	1.0	1	1.0
γ_2	0	1.0	0	2.0
β_1	-1/2	1.0	-1/2	-2.0
β_2	2/3	-1.0	2/3	1.0
b	1.0	0.0	1.0	1.0

Table 4.1: Parameters of the two GIRK integrators for both Strang splitting schemes.

4.3 Results

To test our newly developed GIRK implicit integrators and compare them against the MDIRK integrator, we created a one-dimensional Riemann solver code for gas–dust mixtures called `pigpen`[2]. This is a lightweight and flexible codebase, particularly well-suited for testing the different implicit integrators discussed above. The code supports both an exact Riemann solver and an HLL Riemann solver. However, only the exact solver is used in the tests presented here. Additionally, we employ a simple linear piecewise reconstruction method with a van Leer slope limiter (Van Leer, 1977) to approximate the inter-cell values, before computing the fluxes between neighboring cells. The timestep is selected using a simple CFL condition

$$\Delta t_{\mathrm{CLF}} = C_{\mathrm{CFL}} \frac{\Delta x}{v_{\mathrm{sig}}^{\mathrm{max}}}, \tag{4.40}$$

where $C_{CFL} = 0.1$, Δx is the (constant) spacing between the cells, and the maximum signal velocity is computed as

$$v_{\mathrm{sig}}^{\mathrm{max}} = \max_n \left(c_{\mathrm{s}}^n + v_{\mathrm{g}}^n + \sum_{i=1}^{N_{\mathrm{d}}} v_{\mathrm{d},i}^n \right) \qquad n = 1, ..., N_{\mathrm{cells}} \tag{4.41}$$

where c_{s} is the sound speed and the n index spans over all cells. All the GIRK implicit integrators described in the previous section were implemented in `pigpen`, as well as the MDIRK integrator by Krapp et al. (2024a). These methods will be tested and compared across a suite of standard multi-fluid dust dynamics benchmarks.

4.3.1 DUSTYBOX

In this standard test for dust dynamics, the DUSTYBOX test, we replicate the one-dimensional collision experiment proposed by Huang & Bai (2022). Gas and two distinct dust populations are initialized with different velocities and densities. Under the influence of aerodynamic drag alone, the system evolves toward an equilibrium state at the center-of-mass velocity. The analytical solution to this problem is given by

$$v(t) = v_{\mathrm{COM}} + c_1 e^{\lambda_1 t} + c_2 e^{\lambda_2 t}, \tag{4.42}$$

where v_{COM} is the center-of-mass velocity, and the parameters c_1, c_2, λ_1, and λ_2 are provided in Huang & Bai, 2022, Table 1, along with the corresponding initial conditions. We consider three setups: a non-stiff case (Test A), a stiff case with large drag (Test B), and a stiff case with a large dust-to-gas ratio (Test C). Fig. 4.1 shows the simulation results compared to the analytical solution.
While for Test A we choose the timestep according to the CFL condition in Eq. 4.40, for

[2] `pigpen` can be freely downloaded at [doi:10.5281/zenodo.17329227]https://doi.org/10.5281/zenodo.17329227. In the same repository, all the tests performed in this section can be found as well as a `jupyter` notebook to reproduce the plots

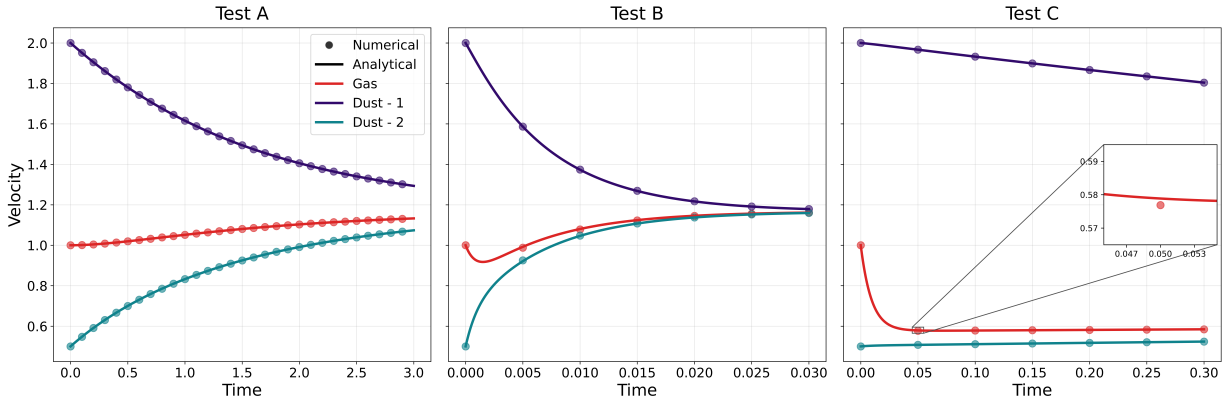

Figure 4.1: Simulation outputs and analytical solutions for the three DUSTYBOX benchmark tests described in Huang & Bai (2022). The system is composed of a gas fluid and two dust fluids. The simulation outputs are shown as dots, while the analytical solution is drawn as a continuous line. For the non-stiff test A, the timestep is chosen according to the CFL condition Eq. 4.40, while for the stiff tests B and C, the timestep is, respectively, $\Delta t = 0.005$ and $\Delta t = 0.05$. The integrator used (GIRK with $\mathcal{D}_{\Delta t/2}\mathcal{H}_{\Delta t}\mathcal{D}_{\Delta t/2}$ Strang splitting) is able to correctly capture the evolution of the mixture for all stiffness regimes. For Test C, we also show how much the integrator undershoots the analytical solution for the gas velocity, before converging.

Tests B and C we fix it to, respectively, $\Delta t = 0.005$ and $\Delta t = 0.05$. This choice follows Krapp et al. (2024a) and is intended to investigate the monotonic convergence of the integrator. Since the GIRK integrator is not fully implicit, it may introduce small oscillations around the expected solution in stiff regimes, thereby preventing strictly monotonic convergence. Fig. 4.1 shows, for Test C, a zoom-in of the gas velocity where a small oscillation is indeed visible before the solution converges. However, the amplitude of this oscillation is significantly smaller than that reported in Fig. 2 of Krapp et al. (2024a).
To quantify accuracy, we compute the relative error of the numerical solution as

$$e = \frac{1}{N_{\text{step}}} \sum_{k=1}^{N_{\text{step}}} \sum_{i=0}^{2} \frac{|v_i(t_k) - v_i^{\text{num}}(t_k)|}{v_i(t_k)}, \tag{4.43}$$

where N_{step} is the number of timesteps, t_k denotes the time at step k, and the index i refers to either the gas phase or one of the two dust phases. By fixing the simulation timestep, we evaluate the relative error for various integration methods on Test A. Fig. 4.2 compares the DIRK (with two drag integration steps) and the GIRK methods (with both two and three drag integration steps, corresponding to the Strang splitting schemes in Eq. 4.20 and 4.38). As shown, the DIRK method achieves second-order accuracy for both large and small timesteps, while the GIRK methods can reach third-order accuracy in some regimes. This behaviour reflects the conditions that we imposed on the GIRK methods in Section 4.2.7.

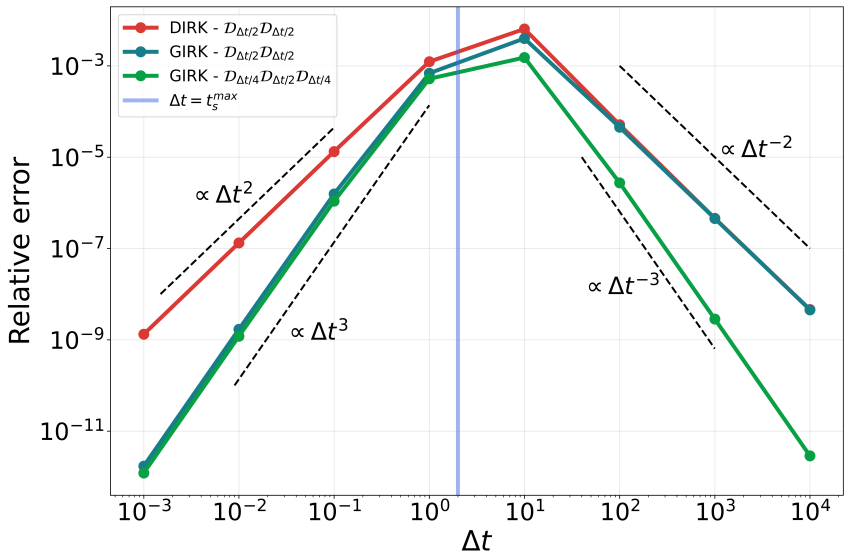

Figure 4.2: Convergence orders for the DIRK and the GIRK integrations for different Strang operator splitting schemes on Test A of Huang & Bai (2022). The error is computed using Eq. 4.43. All integrators follow their expected scalings, given the properties imposed as in Section 4.2.5 and 4.2.7.

4.3.2 External Force Damping

In the next test, we introduce an external force to evaluate the convergence behavior of the integration schemes in the presence of additional source terms. We adopt the benchmark described in Section 3.2 of Krapp et al. (2024a), which provides both the analytical solution and initial conditions. In Fig. 4.3 we show the simulation output for the GIRK integrator in the $\mathcal{D}_{\Delta t/2}\mathcal{H}_{\Delta t}\mathcal{D}_{\Delta t/2}$ splitting scheme. As shown, the numerical solution correctly converges toward the expected asymptotic state. To assess convergence quantitatively, Fig. 4.4 shows the relative error of the final simulation output compared to the analytical asymptotic solution, for various timestep values and integration schemes. Among the tested methods, the MDIRK integrator from Krapp et al. (2024a) delivers the best overall performance. However, it is also the least flexible for general implementation within hydrodynamical codes. The DIRK integrator, when coupled with a Strang splitting scheme, shows, as expected, only first-order convergence for large values of Δt, leading to substantial relative errors when the timestep approaches $t_{\mathrm{s}}^{\mathrm{max}}$. In contrast, the GIRK integrators yield significantly better agreement with the analytical solution and converge more rapidly. While the $\mathcal{D}_{\Delta t/2}\mathcal{H}_{\Delta t}\mathcal{D}_{\Delta t/2}$ splitting scheme delivers second order convergence for $\Delta t > t_{\mathrm{s}}^{\mathrm{max}}$, we notice that the $\mathcal{D}_{\Delta t/4}\mathcal{H}_{\Delta t}\mathcal{D}_{\Delta t/2}\mathcal{H}_{\Delta t}\mathcal{D}_{\Delta t/4}$ offers only a mixed first/second-order convergence.

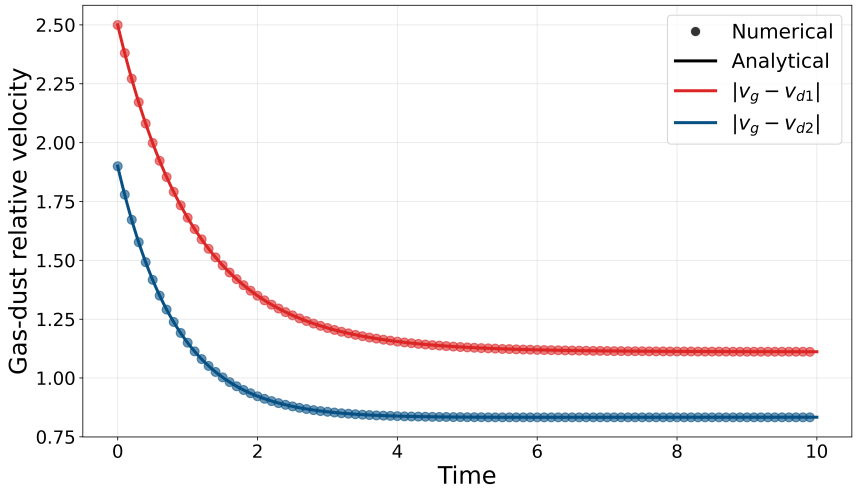

Figure 4.3: External force damping benchmark test. The dots represent the simulation output, while the solid line shows the analytical solution. The simulation was carried out using the GIRK drag integrator with the $\mathcal{D}_{\Delta t/2}\mathcal{H}_{\Delta t}\mathcal{D}_{\Delta t/2}$ Strang splitting scheme. As expected, the simulation follows the expected solution and converges to the correct asymptotic value.

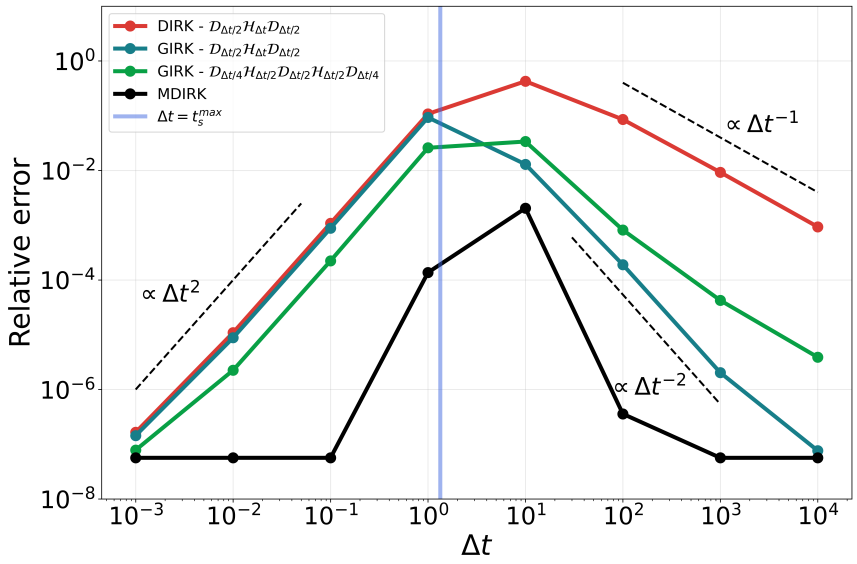

Figure 4.4: Relative error of the final simulation output compared to the analytical asymptotic solution, computed using Eq. 4.43. The MDIRK method performs best, while the DIRK integrator coupled with Strang splitting converges only at first order in the stiff regime. The GIRK integrators, however, show significantly improved accuracy and faster convergence across both splitting schemes.

4.3.3 DUSTYWAVE and Complexity Scaling

The next test is the standard DUSTYWAVE benchmark for the damping of a sound wave, originally presented in Laibe & Price (2012a) and extended to multiple grain sizes

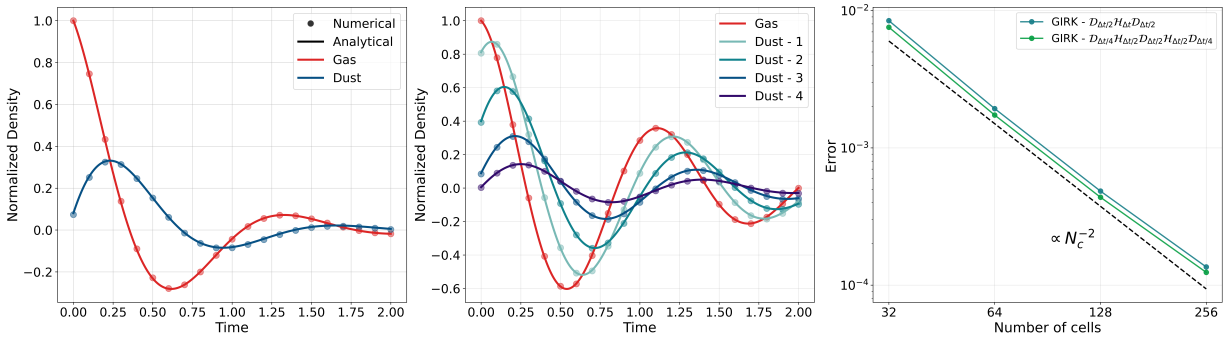

Figure 4.5: Simulation outputs and analytical solution of the DUSTYWAVE benchmark test for both a single dust specie (left panel) and four dust species (central panel). The plots show the time evolution of the normalized density perturbation. The simulation was carried out using the $\mathcal{D}_{\Delta t/2}\,\mathcal{H}_{\Delta t}\,\mathcal{D}_{\Delta t/2}$ Strang splitting scheme. In the right panel, we show the scaling of the error with the number of cells, for both Strang splitting schemes, computed on the test shown in the central panel.

in Benítez-Llambay et al. (2019). This is the first test in which both the hydrodynamical fluxes and the implicit drag terms are evolved simultaneously, making it especially important for assessing the performance of the Strang splitting scheme in coupling hydrodynamics with drag.

The gas and dust properties are initialized as small perturbations of the form

$$\delta f = A\left[\mathrm{Re}(\delta\hat{f})\cos kx - \mathrm{Im}(\delta\hat{f})\sin kx\right], \tag{4.44}$$

where the values for the real and imaginary parts of the perturbations are found in Table 2 of Benítez-Llambay et al. (2019). The amplitude of the perturbation is set to $A = 10^{-4}c_{\rm s}$ for the velocities and $A = 10^{-4}\rho_0$ for the densities. The domain, $L = [0, 1]$, is split into $N = 512$ evenly spaced cells, and the boundary conditions are set to be periodic.

Benítez-Llambay et al. (2019) derived a general solution for any number of dust species, which we use to assess the validity of the numerical solution. Fig. 4.5 shows the agreement between the numerical solution, obtained using the $\mathcal{D}_{\Delta t/2}\,\mathcal{H}_{\Delta t}\,\mathcal{D}_{\Delta t/2}$ Strang splitting scheme, and the analytical solution for both the single dust and the four dust-species cases. In addition, in the right panel of Fig. 4.5, we show the scaling of the error with the number of cells, computed on the test with four dust species. As expected, both Strang splitting methods achieve second-order accuracy with increasing spatial resolution.

Like the MDIRK scheme, the GIRK integrator does not require matrix inversion and is therefore expected to maintain linear complexity as the number of dust species increases. To verify this, we conducted a modified version of the DUSTYWAVE test in which the original dust fluid is artificially split into $N_{\rm d}$ identical fluids, all with the same stopping time but each representing a fraction of the original dust mass. This preserves the analytical solution while increasing the dimensionality of the drag coupling. We then performed simulations for $N_{\rm d}$ ranging from 1 to 512 and measured wall-clock execution time. Fig. 4.6 shows the timing results as a function of $N_{\rm d}$, for the MDIRK integrator and both GIRK

integrators proposed in this work. All methods show tightly clustered results around a linear trend, confirming the expected scaling behavior.

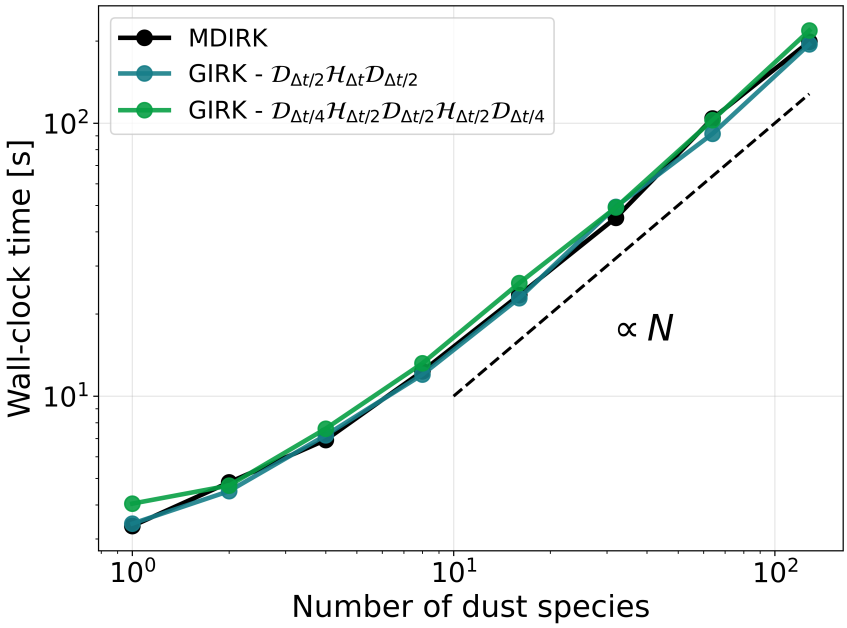

Figure 4.6: Wall-clock execution time for the DUSTYWAVE test for different number of dust species. By using the analytical solution and avoiding matrix inversion, all integrators scale linearly with the number of dust species.

4.3.4 DUSTYSHOCK

Unlike the linear wave dynamics examined in the DUSTYWAVE test, the DUSTYSHOCK benchmark evaluates the performance of drag integrators in a more challenging regime involving shocks and discontinuities. The original shock solution for a single dust species was presented in Lehmann & Wardle (2018), and later generalized to an arbitrary number of dust species by Benítez-Llambay et al. (2019). In this test, we adopt their analytical solution and initial conditions to evaluate the GIRK integrator, simulating a gas mixture with three dust species.

Because the analytical solution describes the post-shock steady state, we compare it with the numerical result only after the shock has fully propagated and the system has relaxed. Following the setup of Benítez-Llambay et al. (2019), we simulate a domain $x \in (0, 40)$, discretized with 400 cells. The left state is imposed on the first 40 cells, while the remaining 360 cells represent the right state, creating a shock initially located at $x = 4$. Transmissive (i.e., zero-gradient) boundary conditions are imposed. All fluids are initialized with the same left and right densities and velocities

$$\rho_i^{\mathrm{L}} = 1.0, \quad \rho_i^{\mathrm{R}} = 16.0, \quad \omega_i^{\mathrm{L}} = 1.0, \quad \omega_i^{\mathrm{R}} = 0.0625, \tag{4.45}$$

where $\omega_i = v_i/(\mathcal{M}c_s)$ is the velocity normalized by the sound speed and the Mach number. The three dust fluids differ only in drag coefficients, being $K_1 = 1.0$, $K_2 = 3.0$ and $K_3 = 5.0$, respectively. The sound speed is set to $c_s = 1.0$ and the Mach number to $\mathcal{M} = 2.0$.
Fig. 4.7 shows the density and normalized velocities of the gas and the three dust species at $t = 500$ after the shock, alongside the numerical solution found by Benítez-Llambay et al. (2019). The integrator used was GIRK with the $\mathcal{D}_{\Delta t/2}\mathcal{H}_{\Delta t}\mathcal{D}_{\Delta t/2}$ Strang splitting scheme. The excellent agreement demonstrates that the GIRK integrator robustly handles drag even in the presence of strong shocks, while maintaining accurate coupling with the hydrodynamics.

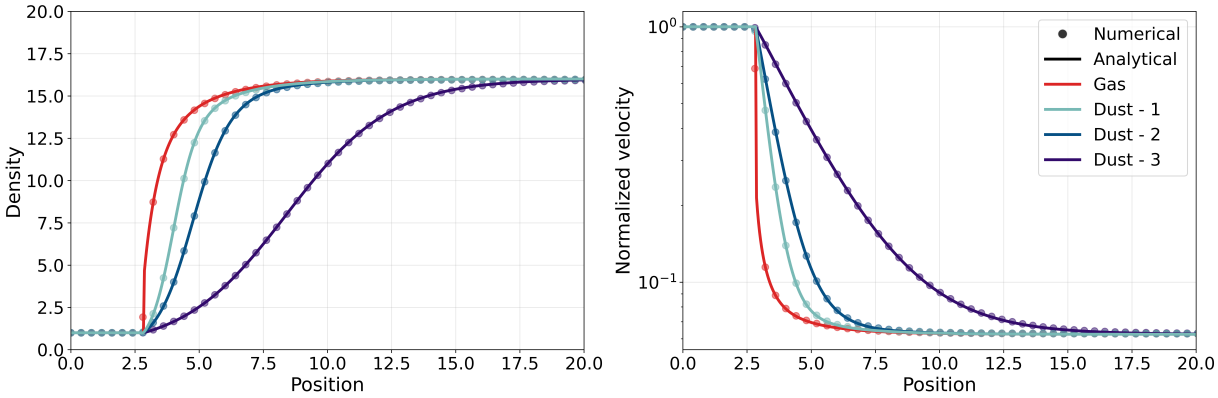

Figure 4.7: Density and normalized velocity of the gas and three dust species for the DUSTYSHOCK benchmark test. The figure shows a snapshot taken at $t = 500$ after the shock to ensure the system has had time to approach its steady state.

4.3.5 Steady-State 1D Shearing Box

We conclude our test suite with the steady-state gas and dust drift problem in a one-dimensional shearing box. This configuration, first introduced by Krapp et al. (2024a), represents the most complex setup in our suite, as it combines hydrodynamical fluxes, external forces, and non-trivial boundary conditions. The system includes centrifugal and Coriolis forces, as well as an additional term that mimics the radial pressure gradient in a protoplanetary disk, acting only on the gas momentum. Periodic boundary conditions are applied, with an extra source term in the y component of the gas and dust momenta to reproduce the background azimuthal shear. The equilibrium solution for this problem was originally derived by Benítez-Llambay et al. (2019) for multiple dust fluids, generalizing the classic single-fluid result of Nakagawa et al. (1986). We adopt the equilibrium configuration described in Section 3.5.2 of Benítez-Llambay et al. (2019) as the initial condition, allowing us to test the ability of our code to maintain this steady state over time.
The GIRK integrator was able, in both Strang splitting schemes, to maintain the initial conditions at equilibrium, with absolute variations with respect to the initial conditions smaller than 10^{-12}, at a time of $t = 20$ in code units. This benchmark confirms that

GIRK is able to resolve the dynamics of dust and gas even in complex environments, with external forces and non-standard boundary conditions.

4.4 Discussion and Conclusion

The integration of the drag term is a critical component in the dynamics of gas–dust mixtures involving multiple dust species. In recent years, several integrators have been developed to address this problem, with a particular emphasis on implicit schemes that avoid the restrictive timestep constraints imposed by small stopping times. While many of these approaches rely on matrix inversions, resulting in cubic computational complexity, the analytical solution introduced by Krapp & Benítez-Llambay (2020) and extended by Krapp et al. (2024a) provides a path to linear-time integration.
In this work, we formalized a systematic procedure to construct drag integration schemes following the analytical framework of Krapp & Benítez-Llambay (2020), allowing us to evaluate their scaling and convergence properties. We first applied this procedure to the DIRK integrator presented in Krapp et al. (2024a), coupled with a Strang splitting scheme, and found that it achieves second-order accuracy across most regimes. However, we also showed that DIRK fails to reach second-order convergence toward the equilibrium solution when external forces are applied and the timestep exceeds the maximum stopping time, i.e., for $\Delta t > t_{\mathrm{s}}^{\max}$. To overcome this limitation, we introduced a novel implicit drag integrator, GIRK (General Implicit Runge–Kutta), which generalizes the DIRK structure by lifting the diagonal constraint and introducing additional tunable parameters. This flexibility enables higher-order convergence while preserving linear complexity. Table 4.2 summarizes the scaling properties of the various methods discussed in this paper, both in terms of accuracy of the drag integrator and convergence in the presence of an external forcing.

	Accuracy		Convergence	
	$\Delta t < t_s^{\max}$	$\Delta t > t_s^{\max}$	$\Delta t < t_s^{\max}$	$\Delta t > t_s^{\max}$
DIRK - $\mathcal{DHD}$	II	II	II	I
GIRK - $\mathcal{DHD}$	III	II	II	II
GIRK - $\mathcal{DHDHD}$	III	III	II	I / II

Table 4.2: Accuracy and convergence orders for the DIRK and GIRK integrators across different timestep regimes.

Compared to the MDIRK integrator proposed in Krapp et al. (2024a), GIRK offers a more versatile implementation into existing hydrodynamical codes through Strang splitting or similar operator-splitting techniques. Moreover, since it is derived using the same analytical foundation as DIRK and MDIRK, GIRK retains their key advantage: linear computational complexity with respect to the number of dust species. This makes it especially suitable for large-scale simulations involving processes like coagulation and fragmentation. Notably,

GIRK relies on the same underlying approximation as DIRK: momentum exchange occurs only between the gas and each dust species, and not among dust species themselves. This assumption may break down in environments with highly concentrated dust clumps, which are uncommon in most astrophysical applications and would require specialized numerical treatments. To further evaluate GIRK's performance, we plan to implement it in more complex hydrodynamical simulation codes, beyond the simple one-dimensional test code `pigpen` used in this work.

Acknowledgments

TB acknowledges funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under grant 325594231. GTP, TB, KD and BE acknowledge support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC-2094 - 390783311. GTP and KD acknowledge support by the COMPLEX project from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program grant agreement ERC-2019-AdG 882679. GTP and TB acknowledge funding from the European Union under the European Union's Horizon Europe Research and Innovation Programme 101124282 (EARLYBIRD). Views and opinions expressed are, however, those of the authors only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. MH acknowledges support from the DFG program "Closing the Loop - Using Synthetic Observations of Simulated Star-forming Regions to Test Observational Properties" (DFG Project Number: 426714422).

Appendix A: Variables for the Full Solution

We report in this appendix the definitions of the constants needed to compute the full analytical solution to the aerodynamic drag for the GIRK integrator to be used in Eq. 4.29 - 4.32.

$$\Lambda_i = \Big(1 + \alpha_i \Delta t \Big(\gamma_1 + \gamma_2 + \alpha_i \Delta t \left(\gamma_1\gamma_2 - \beta_1\beta_2\right)\Big)\Big)^{-1}, \quad \delta_{pi} = \Big(1 + \gamma_p \Delta t \alpha_i\Big)^{-1}, \tag{4.46}$$

$$A_p = \sum_{i=1}^{N_\mathrm{d}} \alpha_i u_{di}\, \delta_{pi} - \beta_p \Delta t \sum_{i}^{N_\mathrm{d}} \alpha_i^2 u_{di} \Big(1 + \alpha_i \Delta t (\gamma_p - \beta_q)\Big) \delta_{pi}\, \Lambda_i, \tag{4.47}$$

$$B_p = \sum_{i}^{N_\mathrm{d}} \alpha_i \epsilon_i\, \delta_{pi} - \beta_p \Delta t \sum_{i}^{N_\mathrm{d}} \alpha_i^2 \epsilon_i \Big(1 + \alpha_i \Delta t \left(\gamma_p - \beta_q\right)\Big) \delta_{pi}\, \Lambda_i, \tag{4.48}$$

$$C_p = \sum_{i}^{N_\mathrm{d}} \alpha_i \epsilon_i\, \delta_{pi} - \Delta t \sum_{i}^{N_\mathrm{d}} \alpha_i^2 \epsilon_i (\gamma_q + \alpha_i \Delta t (\gamma_1\gamma_2 - \beta_1\beta_2)) \delta_{pi} \Lambda_i, \tag{4.49}$$

$$D_p = 1 + \gamma_p \Delta t \sum_i^{N_\mathrm{d}} \alpha_i \epsilon_i \, \delta_{pi} - \beta_1 \beta_2 \Delta t^2 \sum_i^{N_\mathrm{d}} \alpha_i^2 \epsilon_i \delta_{pi} \Lambda. \tag{4.50}$$

The index i spans the N_d dust grain sizes, whereas the index $p = 1, 2$ differentiates between two similar sets of constants, while q takes the opposite value with respect to p.

Appendix B: Stability analysis

In this appendix, we show the stability analysis for both GIRK integrators presented in Section 4.2.7. The `jupyter` notebook used to perform this stability analysis is freely available on `github`[3]

The first step is to introduce the stability function for a general Runge-Kutta method

$$R(\mu) = \frac{\det(\mathbb{1} - \mu \mathbf{A} + \mu(\mathbf{e} \cdot \mathbf{b}))}{\det(\mathbb{1} - \mu \mathbf{A})}, \tag{4.51}$$

where $\mathbf{e} = (1, 1)$ and the $\mathbf{A}$ matrix and $\mathbf{b}$ vector are defined by the Runge-Kutta parameters. Following the choice we made throughout this paper

$$\mathbf{A} = \begin{pmatrix} \gamma_1 & \beta_1 \\ \beta_2 & \gamma_2 \end{pmatrix}, \quad \mathbf{b} = (b, 1 - b). \tag{4.52}$$

The variable $\mu = \lambda \Delta t$ is the product between the timestep Δt and the eigenvalues of the matrix defining the problem at hand; in this case, the matrix $\mathbf{M}$ in Eq. 4.9.

The stability region associated to a Runge-Kutta method is the region in the μ plane where $|R(\mu)| < 1$; consequently, a Runge-Kutta method is considered A-stable if the stability region includes the entire left half-plane of μ.

In Fig. 4.8 we show the stability regions of both Strang splitting schemes introduced in Section 4.2.7. In this context, the only difference between the two methods lies in the choice of parameters, and thus of $\mathbf{A}$ and $\mathbf{b}$. In Krapp & Benítez-Llambay (2020), the inverse of the maximum eigenvalue of $\mathbf{M}$ is derived to be bounded between the two largest stopping times of the system. Hence, the boundary between the two sets of parameters for each method, $\Delta t = t_\mathrm{s}^\mathrm{max}$, corresponds here to the circumference $|\mu| = 1$.

As can be seen in Fig. 4.8, both Strang splitting methods are A-stable, since their stability region contains the entire left half-plane. Although not shown in the Fig., we also checked that both methods are asymptotically stable, or L-stable, by checking that $R(\mu) \to 0$ as $\mu \to \infty$.

[3] https://github.com/Giovanni-Tedeschi/GIRK_stability

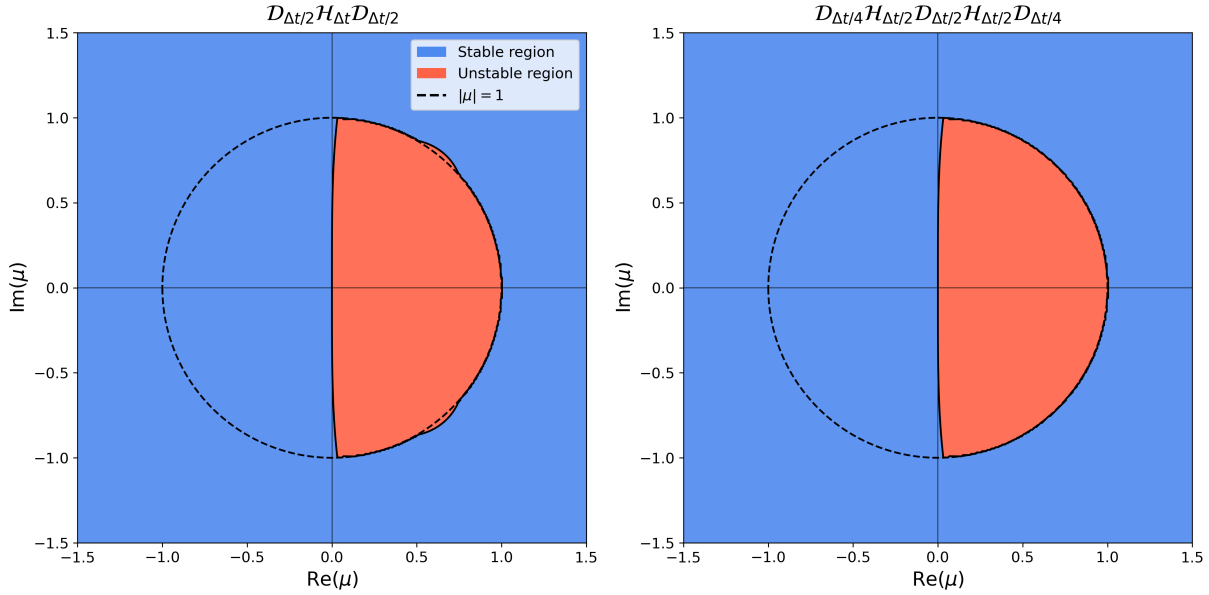

Figure 4.8: Stability regions of both Strang splitting schemes for the GIRK integrator. Being the left half-plane stable, both integrators are A-stable. The $|\mu| = 1$ circumference marks the boundary between the parameters chosen for $\Delta t < t_{\mathrm{s}}^{\mathrm{max}}$ and $\Delta t > t_{\mathrm{s}}^{\mathrm{max}}$.

5

Modeling dust dynamics in `OpenGadget3` II. A Meshless Finite Mass implementation of the multi-grain One–Fluid model

In prep. for submission on Astronomy & Astrophysics as
Tedeschi-Prades et al., 2026

Full authors list:
Giovanni Tedeschi-Prades; Til Birnstiel; Klaus Dolag; Barbara Ercolano; Mark Hutchison; Frederick Groth.

In recent years, a variety of approaches to modeling dust dynamics have been developed within both Lagrangian and Eulerian simulation frameworks. In this paper, we present a new method for simulating dust dynamics by coupling the One–Fluid model to the Meshless Finite Mass (MFM) hydrodynamical framework. In this approach, simulation particles represent a mixture of gas and dust moving at the center-of-mass velocity, while inter-particle fluxes are computed separately for the different phases. The resulting method combines the advantages of the One–Fluid formalism for particle-based simulations with the accuracy and robustness of Riemann-solver-based flux calculations, providing a promising framework for self-consistent gas–dust simulations.
The implementation presented here successfully passes a suite of standard dust-dynamics benchmark tests without requiring artificial viscosity terms, which are commonly employed in SPH implementations of the One–Fluid model. Moreover, the hybrid structure of the method naturally lends itself to future extensions, including magnetohydrodynamic simulations and more advanced dust-dynamics models that can be incorporated within the Riemann-solver framework.

5.1 Introduction

Dust, together with gas, constitutes a fundamental component of a wide range of astrophysical systems, from the diffuse interstellar medium (ISM) and molecular clouds to protoplanetary disks and galactic outflows. In many observational contexts, dust is the most readily detectable constituent, emitting thermally at infrared and millimeter wavelengths, while gas properties are often inferred indirectly under the assumption of a fixed dust-to-gas mass ratio. Accurately modeling the dynamics of dust is therefore essential both for the correct interpretation of observations and for predictive numerical simulations of astrophysical environments.

In many applications, dust is assumed to be perfectly coupled to the gas and treated as a passive scalar advected by the flow. This approximation breaks down when the aerodynamic coupling becomes inefficient, such as for large grain sizes, low gas densities, or elevated dust-to-gas ratios. Protoplanetary disks provide a well-studied example in which differential dust–gas dynamics plays a central role in disk evolution and planet formation (Birnstiel, 2024). In these systems, the radial drift of dust grains, driven by the sub-Keplerian rotation of the gas due to radial pressure support, is a key mechanism governing grain migration, growth barriers, and the redistribution of solids throughout the disk.

The inclusion of self-consistent dust dynamics in hydrodynamical simulation codes has expanded rapidly in recent years. Early approaches often relied on post-processing gas-only simulations with tracer particles, neglecting the dynamical back-reaction of dust on gas. While computationally inexpensive, such methods can lead to inaccurate results whenever dust contributes non-negligibly to the inertia or momentum budget of the system. Fully coupled dust–gas models have since been developed within both Eulerian and Lagrangian frameworks. In grid-based codes, dust is commonly treated either as a pressureless fluid or as a collection of particles coupled to the gas via drag forces (e.g. Huang & Bai, 2022; Benítez-Llambay et al., 2019). In Smoothed Particle Hydrodynamics (SPH), dust dynamics has traditionally been modeled using either two-fluid approaches or the One–Fluid formulation introduced by Laibe & Price (2014b), which describes the dust–gas mixture using a single set of fluid variables.

An alternative class of numerical methods is provided by the Meshless Finite Mass (MFM) and Meshless Finite Volume (MFV) schemes (Lanson & Vila, 2008; Gaburov & Nitadori, 2011; Hopkins, 2015; Groth et al., 2023). These methods combine aspects of Lagrangian particle techniques and Godunov-type finite-volume solvers: resolution elements move with the flow, while fluxes between them are computed by solving Riemann problems. As a result, MFM naturally inherits the advantages of both approaches, including excellent shock-capturing properties, exact conservation of mass, momentum, and energy, and reduced numerical diffusion compared to traditional SPH formulations.

In the previous paper of this series (Tedeschi-Prades et al., 2025b), we described our implementation of the One–Fluid dust model in the SPH framework of `OpenGadget3`, restricted to a single dust species. In the present work, we extend this effort by developing a self-consistent implementation of dust–gas dynamics within the MFM framework

of `OpenGadget3`, coupled to the One–Fluid formalism and generalized to multiple dust species. In this approach, simulation particles represent the mixture of gas and dust, as in the One–Fluid model, while the gas and dust components are treated separately at the level of the Riemann solver when computing inter-particle fluxes. This strategy allows us to combine the advantages of the One–Fluid description with the accurate flux-based nature of MFM.
The structure of the paper is as follows. In Section 5.2, we briefly review the MFM framework and describe how dust dynamics can be incorporated within it using the One–Fluid formalism. Particular attention is devoted to the treatment of dust at the Riemann-solver level and to a new algorithm developed to determine the zero-mass-flux reference frame when computing inter-particle fluxes. In Section 5.3, we validate the implementation using a suite of standard benchmark problems. Finally, in Section 5.4, we summarize our results and discuss possible extensions and applications of the method.

5.2 Methods

5.2.1 Meshless Finite Mass in `OpenGadget3`

This section briefly reviews the implementation of the Meshless Finite Mass (MFM) method in `OpenGadget3`. A complete and detailed description of the numerical scheme and its implementation can be found in Groth et al. (2023). Here, we summarize the key ideas and notation required for the discussion of dust dynamics in the following sections.
A hydrodynamical system governed by conservation laws can generally be written as a system of partial differential equations of the form

$$\frac{\partial \vec{U}}{\partial t} + \nabla \cdot \left(\vec{F}(\vec{U}) - \vec{v}_\mathrm{f} \otimes \vec{U} \right) = 0, \tag{5.1}$$

where $\vec{U}$ denotes the vector of conserved variables, $\vec{F}(\vec{U})$ is the corresponding flux vector, and $\vec{v}_\mathrm{f}$ is the velocity of the reference frame in which the equations are written. The symbol $\otimes$ denotes the outer product. For simplicity, we neglect here the presence of source terms, which would otherwise appear on the right-hand side of Eq. (5.1).
Well-known examples of systems that can be cast in the form of Eq. (5.1) include the Euler equations of gas dynamics and the equations of ideal magnetohydrodynamics (MHD). In contrast, systems such as the Navier–Stokes equations or non-ideal MHD introduce additional non-conservative terms (e.g. viscosity, resistivity), whose fluxes cannot be written solely as functions of $\vec{U}$ and therefore require separate numerical treatments.
Following the work of Lanson & Vila (2008), Gaburov & Nitadori (2011), and Hopkins (2015), a conservative system such as Eq. (5.1) can be discretized within a mesh-free framework by associating the conserved quantities with a set of moving particles. In this approach, the discretized form of Eq. 5.1 reads

$$\frac{\mathrm{d}}{\mathrm{d}t} \left(V_i \, \vec{U}_i \right) + \sum_{j \in \mathrm{Ngb}(i)} \left(\tilde{\vec{F}}_{ij} \cdot \vec{A}_{ij} \right) = 0, \tag{5.2}$$

where V_i is the effective volume associated with particle i, $\vec{U}_i$ is the vector of conserved variables carried by that particle. The vector $\vec{A}_{ij}$ represents the effective interface area between particles i and j, while $\tilde{\vec{F}}_{ij}$ is the numerical flux, and the sum extends over its neighboring particles.
This formulation closely mirrors that of a finite-volume method, with the crucial difference that the mesh is replaced by a set of moving, mesh-free resolution elements. As a result, MFM retains the conservative and shock-capturing properties of Godunov-type schemes, while naturally adapting to the Lagrangian motion of the flow and avoiding mesh distortion issues that typically arise in grid-based approaches.
We briefly review here the sequence of operations required to compute the inter-particle fluxes $\tilde{\vec{F}}_{ij}$ for a given pair of particles in the MFM scheme.

- Compute the velocity of the effective moving interface, $\vec{v}_{\mathrm{f},ij}$, using a second-order accurate estimate
$$\vec{v}_{\mathrm{f},ij} = s_j \vec{v}_j + s_i \vec{v}_i, \tag{5.3}$$
where $\vec{v}_i$ and $\vec{v}_j$ are the velocities of particles i and j, respectively. The weights s_i and s_j are defined as
$$s_i = \frac{h_i}{h_i + h_j}, \qquad s_j = \frac{h_j}{h_i + h_j}, \tag{5.4}$$
with h_i and h_j the corresponding smoothing lengths. This choice ensures symmetry and second-order accuracy in the interface motion.

- Transform the conserved variables $\vec{U}_i$ and $\vec{U}_j$ into the reference frame of the moving interface by applying a Galilean boost with velocity $\vec{v}_{\mathrm{f},ij}$. Working in this frame reduces numerical diffusion and improves the accuracy of the Riemann solution for moving fluids.

- Rotate the boosted states $\vec{U}_i$ and $\vec{U}_j$ such that the separation vector
$$\vec{r}_{ij} = \vec{x}_j - \vec{x}_i \tag{5.5}$$
lies along the x-axis. This step reduces the multidimensional Riemann problem to an effectively one-dimensional problem along the line connecting the two particles.

- Reconstruct the left and right interface states of $\vec{U}_i$ and $\vec{U}_j$ at the particle interface using a suitable spatial reconstruction scheme. This step typically involves slope limiting to preserve monotonicity and prevent spurious oscillations near discontinuities.

- Solve the resulting one-dimensional Riemann problem using the reconstructed left and right states. The choice of Riemann solver depends on the specific system of equations being evolved and on the desired balance between accuracy, robustness, and computational cost.

- Transform the resulting numerical fluxes back from the interface frame to the simulation (laboratory) frame by applying the inverse Galilean boost.

Once all these steps have been carried out for each interacting particle pair, the fluxes can be inserted into Eq. (5.2) to compute the time derivatives of the conserved variables. Combined with an appropriate time-integration scheme, this procedure advances the solution to the next timestep while maintaining conservation and second-order accuracy.

5.2.2 The One-Fluid Model

To include dust dynamics in the MFM module of `OpenGadget3`, our starting point is the previous paper of this series, Tedeschi-Prades et al. (2025b). In that work, we described the implementation of the One–Fluid model introduced by Laibe & Price (2014b) and Laibe & Price (2014c) within `OpenGadget3`. In this framework, simulation particles represent a mixture of gas and dust that evolves as a single fluid moving with the center-of-mass velocity of the mixture.
In contrast to Tedeschi-Prades et al. (2025b), where only a single dust population was considered, we extend the implementation here to account for multiple dust grain sizes when solving the dust dynamics within the MFM framework. This generalization is essential for capturing size-dependent drag forces, differential settling, and grain-size-dependent transport, all of which play a crucial role in realistic astrophysical environments.
The One–Fluid variables carried by the simulation particles are defined as

$$\rho = \rho_{\rm g} + \sum_k \rho_{{\rm d},k}, \tag{5.6}$$

$$\epsilon_k = \frac{\rho_{{\rm d},k}}{\rho}, \tag{5.7}$$

$$\vec{v} = \frac{\rho_{\rm g}\vec{v}_{\rm g} + \sum_k \rho_{{\rm d},k}\vec{v}_{{\rm d},k}}{\rho}, \tag{5.8}$$

$$\vec{\Delta v}_k = \vec{v}_{{\rm d},k} - \vec{v}_{\rm g}, \tag{5.9}$$

where the index k labels the dust species, corresponding to different grain sizes, and the summation runs over all dust species present in the mixture.
It is also useful to define the total dust fraction and the dust-fraction-weighted relative velocity,

$$\epsilon = \sum_k \epsilon_k, \tag{5.10}$$

$$\vec{\Delta v} = \frac{1}{\epsilon}\sum_k \epsilon_k \vec{\Delta v}_k, \tag{5.11}$$

which allow the gas and dust properties to be expressed compactly in terms of One–Fluid variables:

$$\rho_{\rm g} = \rho(1-\epsilon), \tag{5.12}$$

$$\vec{v}_{\rm g} = \vec{v} - \epsilon\vec{\Delta v}, \tag{5.13}$$

$$\rho_{{\rm d},k} = \rho\epsilon_k, \tag{5.14}$$

$$\vec{v}_{{\rm d},k} = \vec{v} - \epsilon\vec{\Delta v} + \vec{\Delta v}_k. \tag{5.15}$$

Using these definitions, the Euler equations for gas and the pressureless fluid equations for dust can be recast into a conservative system of partial differential equations of the form given in Eq. 5.2, following (Laibe & Price, 2014a). The corresponding vectors of conserved variables and fluxes are

$$\vec{U} = \begin{pmatrix} \rho \\ \rho\epsilon_k \\ \rho\vec{v} \\ \rho\epsilon_k(\vec{v} - \epsilon\vec{\Delta v} + \vec{\Delta v}_k) \end{pmatrix}, \tag{5.16}$$

$$\vec{F} = \begin{pmatrix} \rho\vec{v} \\ \rho\epsilon_k(\vec{v} - \epsilon\vec{\Delta v} + \vec{\Delta v}_k) \\ \rho\vec{v}\otimes\vec{v} + P\mathbb{1} + \rho\sum_k \left[\epsilon_k\vec{\Delta v}_k\left(\vec{\Delta v}_k - \epsilon\vec{\Delta v}\right)\right] \\ \rho\epsilon_k(\vec{v} - \epsilon\vec{\Delta v} + \vec{\Delta v}_k)\otimes(\vec{v} - \epsilon\vec{\Delta v} + \vec{\Delta v}_k) \end{pmatrix}. \tag{5.17}$$

The resulting system is formally conservative and can therefore be discretized as in Eq. 5.2 and evolved within the MFM framework. However, an important complication arises from the mathematical nature of the system: it is not strictly hyperbolic. This can be seen most clearly by isolating the second and fourth equations, which govern the evolution of the dust densities and momenta, and reduce to the pressureless fluid equations

$$\partial_t\rho_{\mathrm{d},k} + \nabla\cdot(\rho_{\mathrm{d},k}\vec{v}_{\mathrm{d},k}) = 0, \tag{5.18}$$

$$\partial_t(\rho_{\mathrm{d},k}\vec{v}_{\mathrm{d},k}) + \nabla\cdot(\rho_{\mathrm{d},k}\vec{v}_{\mathrm{d},k}\otimes\vec{v}_{\mathrm{d},k}) = 0. \tag{5.19}$$

Pressureless systems are only weakly hyperbolic and admit solutions such as multi-streaming and the formation of δ-shocks, which cannot be captured by standard Riemann solvers designed for strictly hyperbolic systems. As a consequence, it is not possible to solve the full One–Fluid system using a single conventional Riemann solver.
Instead, we split the gas and dust components at the level of the Riemann solver, while retaining the One–Fluid formulation for the overall structure of the simulation. This strategy allows us to:

1. employ any standard Riemann solver for the gas component, including straightforward extensions to magnetohydrodynamics (MHD), and
2. use specialized Riemann solvers for the dust component that are specifically designed to handle the weakly hyperbolic nature of pressureless flows, without contaminating the gas solution.

5.2.3 Dust Riemann Solver

For the dust component of the equations, we adopt the Riemann solver introduced by Huang & Bai (2022), which is specifically designed to handle pressureless systems. The

numerical flux across the interface between particles i and j is defined as

$$F_{ij} = \begin{cases} 0, & u_i < 0 \text{ and } u_j > 0, \\ F_i, & u_i > 0 \text{ and } u_j > 0, \\ F_j, & u_i < 0 \text{ and } u_j < 0, \\ F_i + F_j, & u_i > 0 \text{ and } u_j < 0. \end{cases} \tag{5.20}$$

Here, particles i and j represent the left and right states of the Riemann problem, respectively, and u denotes the velocity component along the x-direction. As customary in meshless and finite-volume methods, the state vectors are first rotated so that the inter-particle separation vector lies along the x-axis, reducing the Riemann problem to one spatial dimension.
This solver is tailored to pressureless flows, where the characteristic structure is degenerate and the system admits multi-streaming solutions. In particular, the case $u_i < 0$ and $u_j > 0$ corresponds to particles moving away from the interface, resulting in a vanishing flux, while the case $u_i > 0$ and $u_j < 0$ represents converging streams, for which the fluxes from both sides are summed. This behavior effectively captures the formation of dust concentrations and δ-shocks, which are characteristic features of pressureless dynamics.

5.2.4 Zero-Mass-Flux reference frame finder

One of the defining features of the Meshless Finite Mass (MFM) method is that neighboring particles exchange fluxes in a reference frame chosen such that there is no net mass transfer between them, i.e.,

$$F_\rho = 0.$$

In the gas-only implementation of MFM, this condition can be enforced straightforwardly by moving the reference frame to the contact discontinuity between the left and right states of the Riemann problem. In this frame, the mass flux vanishes by construction, ensuring that particle masses remain strictly constant in time.
In the multifluid case, however, identifying an appropriate zero-mass-flux reference frame becomes significantly more complex. One possible strategy would be to enforce a zero mass flux condition independently for each fluid component. This approach, however, suffers from several fundamental drawbacks. First, it would require solving the Riemann problem for each fluid in a different reference frame, which would break the Galilean invariance of the numerical scheme and lead to inconsistencies when combining fluxes from different components. Second, one of the key motivations behind the One–Fluid formalism is that dust is allowed to move between particles via the evolution of the dust fraction, rather than being tied to the mass of individual resolution elements. Enforcing a zero-mass-flux condition for the dust component would explicitly prevent this exchange and would therefore undermine one of the core advantages of the One–Fluid model.
To preserve both Galilean invariance and the ability of dust to flow across particles, we

instead enforce a zero total mass flux condition for the mixture. Specifically, we seek a reference frame velocity v_{ZMF} such that the sum of the gas and dust mass fluxes vanishes,

$$F_{\rho_{\mathrm{g}}} + \sum_k F_{\rho_{\mathrm{d},k}} = 0, \tag{5.21}$$

where the sum is taken over all dust species. In this frame, individual components are allowed to exchange mass across particle interfaces, but the total mass of each particle remains constant, preserving the defining property of the MFM method.
We have developed a dedicated algorithm to determine the zero-mass-flux reference frame velocity v_{ZMF} in the presence of multiple dust species. The details of this procedure are described in Appendix 5.4.

5.2.5 Drag integrator

The final component of the implementation concerns the integration of aerodynamic drag between the gas and dust fluids. To this end, we adopt the General Implicit Runge–Kutta (GIRK) integrator introduced in Tedeschi-Prades et al. (2025a), which generalizes the implicit scheme of Krapp et al. (2024b) to arbitrary Strang-splitting time integration frameworks. GIRK is second-order accurate in time and converges to the correct equilibrium solution with second-order accuracy in the presence of external forces. Importantly, its computational cost scales linearly with the number of dust species, making it particularly well suited for simulations that include a large number of grain sizes.
These properties are especially desirable in the context of gas–dust dynamics, where the drag force can become stiff for strongly coupled grains and where explicit integration schemes would impose prohibitively small timesteps. By treating the drag term implicitly, GIRK allows stable and accurate integration across a wide range of stopping times, from tightly coupled to weakly coupled regimes, without introducing additional timestep constraints beyond those already required by the hydrodynamics.
The GIRK integrator is incorporated into the global timestepping scheme of `OpenGadget3` using a Strang-splitting approach of the form

$$\vec{U}^{n+1} = \mathcal{D}_{\Delta t/2}\, \mathcal{H}_{\Delta t}\, \mathcal{D}_{\Delta t/2}\, \vec{U}^{n}, \tag{5.22}$$

where the operator $\mathcal{D}$ represents the drag update performed by the GIRK integrator, and $\mathcal{H}$ denotes the hydrodynamical update, implemented through the full Kick–Drift–Kick hydrodynamical step of `OpenGadget3`.

5.3 Tests

5.3.1 DUSTYBOX

The first benchmark tests we performed are designed to validate the implementation of the aerodynamic drag integrator. In particular, we consider the DUSTYBOX test, which

consists of a one-dimensional domain filled with gas and dust fluids with homogeneous densities and velocities. Under these conditions, no hydrodynamical fluxes are generated between particles, and the system evolution is governed exclusively by the mutual aerodynamic drag between gas and dust. As a result, the relative velocity between the fluids decays exponentially toward equilibrium.
Huang & Bai (2022) provide a set of initial conditions and analytical solutions for the DUSTYBOX benchmark, considering a gas phase coupled to two dust fluids. The general form of the analytical solution for the velocity of each phase can be written as

$$v(t) = v_{\mathrm{COM}} + c_1 e^{\lambda_1 t} + c_2 e^{\lambda_2 t}, \tag{5.23}$$

where v_{COM} is the center-of-mass velocity of the system. The coefficients c_1, c_2 and the eigenvalues λ_1, λ_2 depend on the drag coefficients and dust fractions, and are reported for each fluid in Table 1 of Huang & Bai (2022).
We initialize $N = 100$ particles uniformly distributed in the domain $x \in [0, 1]$, adopting periodic boundary conditions. Huang & Bai (2022) define three distinct test configurations. The initial velocities are identical in all cases, but the tests differ in their drag coefficients and dust fractions. Test A corresponds to a non-stiff regime, characterized by relatively small drag coefficients and low dust fractions. Test B is stiff due to large drag coefficients, which cause the dust velocities to equilibrate rapidly with the gas. Finally, Test C is stiff because of the large dust fractions in both dust fluids, resulting in a rapid relaxation of the gas velocity toward the dust-dominated center-of-mass velocity.
The results of these tests are shown in Fig. 5.1. In all cases, the numerical solutions closely match the analytical predictions. This excellent agreement demonstrates that the GIRK integrator accurately captures the drag-driven momentum exchange between gas and dust, and remains stable and accurate across both non-stiff and stiff coupling regimes.

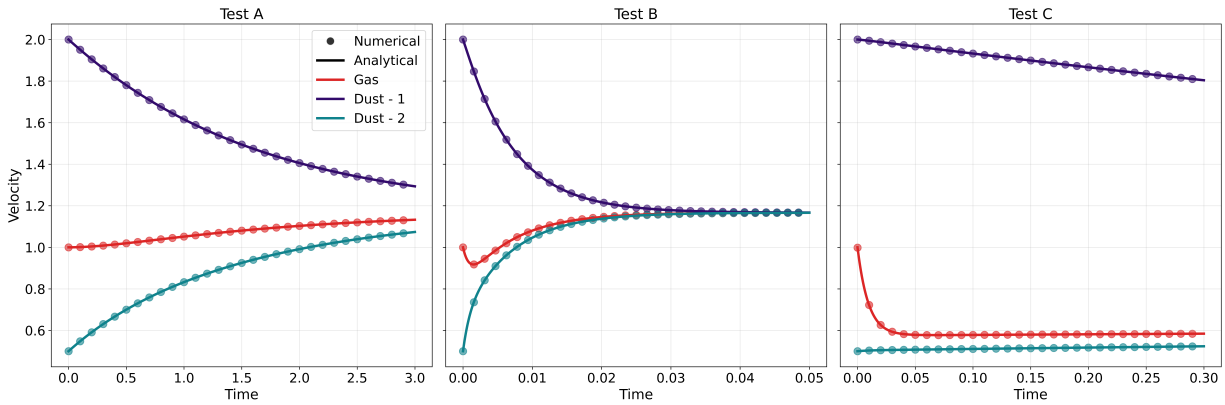

Figure 5.1: Numerical results compared with the analytical solutions for the DUSTYBOX benchmark test described in Huang & Bai (2022). The GIRK integrator accurately reproduces the time evolution of the relative velocities between the gas and the two dust fluids.

5.3.2 DUSTYWAVE

The first benchmark that tests the performance of the full numerical framework, including both hydrodynamical fluxes and aerodynamic drag, is the DUSTYWAVE test. Originally introduced by Laibe & Price (2012a) and later extended to multiple dust species by Benítez-Llambay et al. (2019), this benchmark is specifically designed to assess the correct coupling between gas and dust through drag, as well as its integration within the global timestepping scheme via Strang splitting. As such, it provides a stringent validation of the combined hydrodynamical and drag solvers implemented in `OpenGadget3`.
We initialize $N = 64$ particles uniformly distributed in the one-dimensional domain $x \in [0, 1]$ with periodic boundary conditions. The particle positions are then perturbed to reproduce the density field prescribed in Benítez-Llambay et al. (2019). It is important to note that the density field imposed at initialization corresponds to the One–Fluid formulation, i.e. it represents the total density of the gas–dust mixture. The individual density perturbations for the gas and the dust phases are subsequently recovered by assigning appropriate dust fractions to the particles, consistent with the One–Fluid variables.
Fig. 5.2 shows the numerical solutions compared with the analytical predictions for the normalized densities of gas and dust. Results are presented for both the single dust species case (left panel) and the multi-species case with four distinct dust fluids (right panel). In all cases, the simulations accurately reproduce the analytical wave solutions for each component. This demonstrates that the GIRK drag integrator, when coupled to the MFM hydrodynamics via Strang splitting, correctly captures the joint evolution of gas and dust and remains stable and accurate even in the presence of multiple dust species with different stopping times.

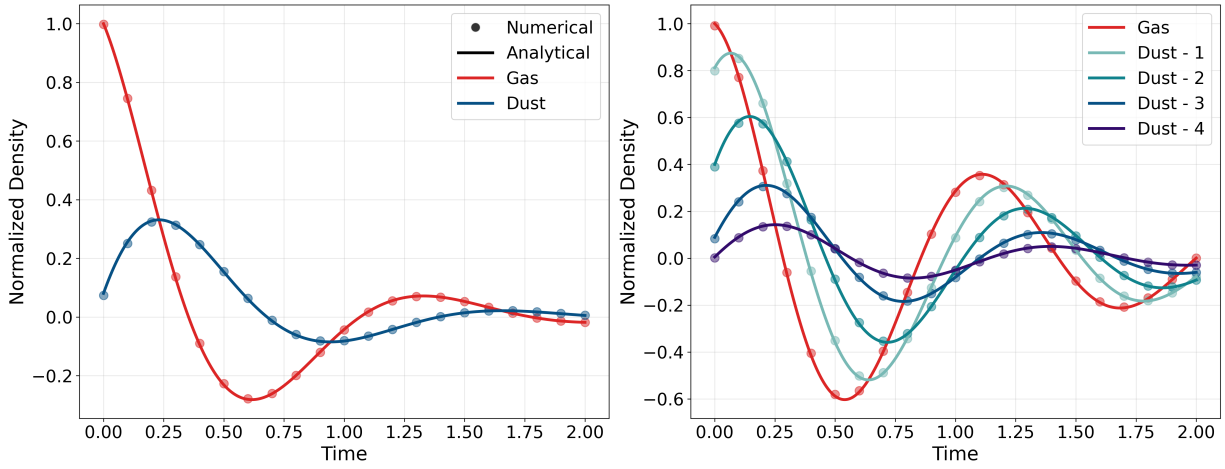

Figure 5.2: Numerical results and analytical solutions for the DUSTYWAVE benchmark test from Benítez-Llambay et al. (2019) for both a single dust fluid (left panel) and four dust fluids (right panel). The agreement between numical ouputs and analytical solutions shows that the GIRK integrator is correctly coupled to the hydrodynamical time stepping of `OpenGadget3` via Strang splitting.

5.3.3 DUSTYSHOCK

The DUSTYSHOCK benchmark test was introduced by Laibe & Price (2012a) to assess how dust responds to the dynamics of an underlying gas flow undergoing a shock. The test is typically performed in two limiting regimes. In the first, dust and gas are uncoupled ($K = 0$), in which case the dust is expected to remain stationary while the gas evolves through a shock. In the second, dust and gas are tightly coupled ($K \gg 1$), so that dust is efficiently dragged by the gas and the mixture behaves as a single fluid.
In the tightly coupled case, when the dust fraction is large, the dynamics of the system are dominated by the inertia of the mixture rather than by the gas alone. As a result, the shock propagates as a single-fluid wave with modified dynamical properties compared to a pure gas shock. In both limiting cases, fiducial solutions can be obtained using standard gas-only Riemann solvers, providing a convenient reference for validating numerical implementations.
Compared to Tedeschi-Prades et al. (2025b), here we combine both regimes within a single simulation by including two dust fluids simultaneously. The left and right initial densities for gas and dust are set to

$$\begin{aligned} \rho_L^{\mathrm{g}} &= 1, \qquad \rho_R^{\mathrm{g}} = 0.125, \\ \rho_L^{\mathrm{d}_1} &= 1, \qquad \rho_R^{\mathrm{d}_1} = 0.125, \\ \rho_L^{\mathrm{d}_2} &= 0.125, \qquad \rho_R^{\mathrm{d}_2} = 0.125, \end{aligned}$$

where the first dust fluid is tightly coupled to the gas ($K_1 = 10^6$), while the second dust fluid is completely uncoupled ($K_2 = 0$). All fluids are initially at rest, and the gas pressure is initialized as $P_L = 1$ and $P_R = 0.1$.
The simulation employs $N = 500$ particles in the domain $x \in [0, 1]$, distributed between the left and right states according to the total density of the mixture. Reflective boundary conditions are imposed via boundary particles placed at the edges of the domain.
The results are shown in Fig. 5.3, which presents the evolution of density, velocity, and internal energy for the gas and both dust fluids. The simulation correctly captures the expected behavior in both coupling regimes: the tightly coupled dust fluid is swept along by the gas shock and evolves coherently with it, while the uncoupled dust fluid remains stationary, unaffected by the gas dynamics.
Notably, in contrast to the SPH implementation presented in Tedeschi-Prades et al. (2025b), no additional artificial viscosity terms are required to recover the correct solution. In particular, the SPH version of the One–Fluid model struggles to maintain the uncoupled ($K_2 = 0$) dust fluid at rest without introducing non-conservative, ad hoc dissipation terms. The MFM-based implementation presented here naturally avoids this issue, highlighting one of the advantages of the meshless finite-mass framework for multifluid dust–gas dynamics.

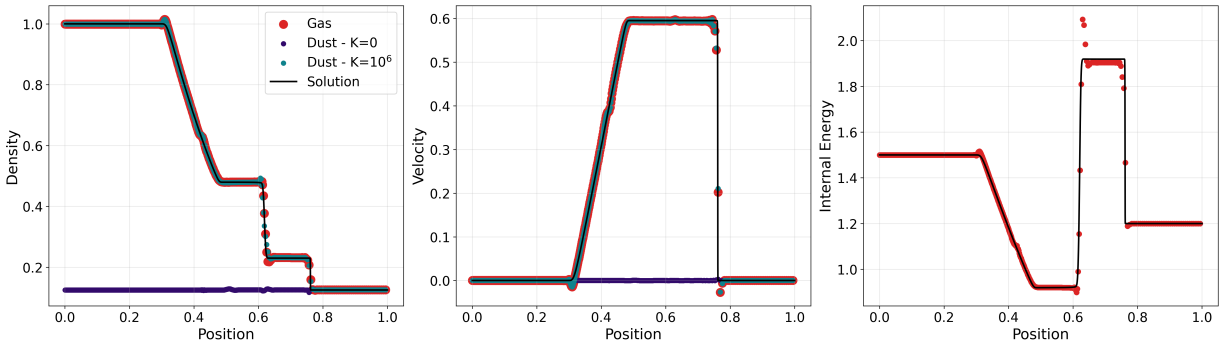

Figure 5.3: Results of the DUSTYSHOCK benchmark test. Contrary to Tedeschi-Prades et al. (2025b), we include here both the uncoupled and tightly coupled dust fluids in the same simulation. The meshless-finite-mass framework is able to recover the correct behaviour of the uncoupled dust fluid without the need for specialized artificial viscosity, contrary to the SPH implementation of the One–Fluid model.

5.4 Conclusion

In this paper, we have presented a new method for solving dust dynamics within the One–Fluid formalism using the MFM hydrodynamical framework. In particular, we have discussed:

1. how the equations governing dust dynamics can be incorporated into the MFM framework by exploiting their conservative formulation;

2. why it is both valid and advantageous to split the One–Fluid variables into separate gas and dust Riemann problems, and to combine their solutions when advancing the system in time;

3. a new algorithm to determine the zero-mass-flux reference frame, which ensures that particle masses remain constant while still allowing dust mass to be exchanged between particles;

4. how aerodynamic drag is integrated using the GIRK algorithm of Tedeschi-Prades et al. (2025a), and how this integrator is coupled to the hydrodynamical update of `OpenGadget3` through a Strang-splitting scheme.

We validated our implementation against the standard DUSTYBOX, DUSTYWAVE, and DUSTYSHOCK benchmark tests, which jointly probe the correct treatment of drag, numerical fluxes, and their coupling through operator splitting. In all cases, the numerical results are in excellent agreement with the corresponding analytical or reference solutions. Notably, the MFM-based implementation successfully reproduces all benchmarks without the need for additional artificial viscosity terms, which were required in the SPH-based implementation presented in (Tedeschi-Prades et al., 2025b).

An important advantage of the present approach is the explicit splitting between gas and dust at the level of the Riemann solver. This design choice naturally enables a number of straightforward extensions. On the gas side, the method can be directly coupled to the magnetohydrodynamic (MHD) version of the MFM solver, allowing the study of dust dynamics in magnetized flows without modification of the dust treatment. On the dust side, the Riemann solver can be readily replaced or extended to incorporate additional physical effects, such as grain charging or the treatment of particle-crossing trajectories, without impacting the gas solver. These features make the proposed framework a flexible and robust foundation for future studies of dust–gas dynamics in a wide range of astrophysical environments.

Acknowledgements

TB acknowledges funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under grant 325594231. GTP, TB, KD, and BE acknowledge support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC-2094 - 390783311. GTP and KD acknowledge support by the COMPLEX project from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program grant agreement ERC-2019-AdG 882679. GTP and TB acknowledge funding from the European Union under the European Union's Horizon Europe Research and Innovation Programme 101124282 (EARLYBIRD). Views and opinions expressed are, however, those of the authors only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. MH acknowledges support from the DFG program "Closing the Loop - Using Synthetic Observations of Simulated Star-forming Regions to Test Observational Properties" (DFG Project Number: 426714422).

Appendix A: Zero-Mass-Flux reference frame finder

In this appendix we describe the algorithm developed to determine the zero-mass-flux (ZMF) reference frame in the presence of multiple fluids. In the present work, the algorithm is applied to a mixture of gas and multiple dust species, but it is fully general and can be employed in any multifluid context, provided that the density flux of each fluid can be evaluated in an arbitrary reference frame.
The goal of the algorithm is to find the reference-frame velocity $v_{\rm ZMF}$ such that the total density flux across the interface vanishes,

$$F_\rho(v_{\rm ZMF}) = 0, \tag{5.24}$$

where

$$F_\rho(v) \equiv F_{\rho_{\rm g}}(v) + \sum_k F_{\rho_{{\rm d},k}}(v) \tag{5.25}$$

is the total mass flux of gas and all dust species evaluated in a frame moving with velocity v along the Riemann problem direction.
We begin by evaluating the total density flux for an initial guess of the reference-frame velocity, v^0_{ZMF}, yielding $F_\rho(v^0_{\mathrm{ZMF}})$. If the absolute value of this flux is already smaller than a prescribed tolerance δF_ρ, the algorithm terminates and returns v^0_{ZMF}.
If this condition is not satisfied, the algorithm enters an iterative procedure aimed at refining the estimate of v_{ZMF}. The method is inspired by a Newton–Raphson root-finding scheme, but avoids the explicit computation of derivatives by using finite differences. At each iteration, we evaluate the flux at a perturbed velocity,

$$F_\rho(v^0_{\mathrm{ZMF}} + \delta v_{\mathrm{ZMF}}), \tag{5.26}$$

where δv_{ZMF} is a small velocity increment.
Using these two flux evaluations, the reference-frame velocity is updated according to

$$v^1_{\mathrm{ZMF}} = v^0_{\mathrm{ZMF}} + \frac{F_\rho(v^0_{\mathrm{ZMF}} + \delta v_{\mathrm{ZMF}})}{F_\rho(v^0_{\mathrm{ZMF}} + \delta v_{\mathrm{ZMF}}) + F_\rho(v^0_{\mathrm{ZMF}})}\,\delta v_{\mathrm{ZMF}}. \tag{5.27}$$

The total density flux is then re-evaluated at the updated velocity v^1_{ZMF}. If

$$\left|F_\rho(v^1_{\mathrm{ZMF}})\right| < \delta F_\rho, \tag{5.28}$$

the algorithm is considered converged and v^1_{ZMF} is returned. Otherwise, the same procedure is repeated iteratively, updating the reference-frame velocity until convergence is achieved or a maximum number of iterations $i_{\max}$ is reached.
In our implementation, we adopt the following numerical parameters:

$$v^0_{\mathrm{ZMF}} = 0, \qquad \delta F_\rho = 10^{-6}, \qquad \delta v_{\mathrm{ZMF}} = 10^{-12}, \qquad i_{\max} = 20. \tag{5.29}$$

These values were found to provide robust and rapid convergence in all tested configurations.
The resulting ZMF velocity ensures that the MFM discretization remains strictly conservative with respect to total mass, while preserving Galilean invariance and the ability of dust fractions to move across particles, which are essential properties of the One–Fluid formulation.

6

Summary and Future Perspectives

Summary of the Thesis

The focus of this thesis has been the development of new numerical methods to simulate the dynamics of dust in a self-consistent manner, without relying on limiting approximations that restrict the accessible parameter space to dust grains tightly coupled to the gas. The overarching goal was, and remains, to enable hydrodynamical simulations to explore regimes in which dust grains exhibit dynamics that differ significantly from those of the gas, with the longer-term objective of assessing the impact of this extended parameter space on simulations of astrophysical systems.
In particular, this thesis has achieved the following:

1. I implemented the One–Fluid model for dust dynamics within the SPH framework of `OpenGadget3`, and improved upon existing implementations in several key aspects. First, I derived and implemented an SPH discretization of the corrected One–Fluid equations without restricting the treatment to the terminal velocity approximation. When exploring the dynamics of extremely weakly coupled dust fluids, I encountered numerical instabilities that motivated the development of a diffusion model capable of regularizing the solution in these extreme regimes, while leaving smooth flows essentially unaffected. This work highlighted a fundamental limitation shared by all approaches that treat dust as a fluid, independent of the underlying hydrodynamical framework: the inability to represent interpenetrating particle trajectories. Introducing diffusion, whether physically motivated or purely numerical, provides a practical and robust way to mitigate this limitation and maintain numerical stability.

2. Prior to extending the model to multiple grain sizes, I identified the opportunity to improve upon the MDIRK implicit drag integrator of Krapp et al. (2024b) by generalizing it to arbitrary Strang-splitting schemes. While the original MDIRK approach

is powerful, its reliance on a specialized timestepping strategy limits its applicability in existing hydrodynamical codes. The resulting General Implicit Runge–Kutta (GIRK) integrator developed in this thesis is fully compatible with Strang splitting, making it straightforward to incorporate into a wide range of simulation frameworks. This generalization enables simulations with a large number of dust species at a computational cost that scales linearly with the number of grain sizes, a crucial feature for realistic dust evolution models.

3. Building upon the GIRK integrator and the One–Fluid infrastructure implemented in `OpenGadget3`, I extended the treatment of dust dynamics to multiple dust species within the Meshless Finite Mass (MFM) module. This step introduced additional challenges, as the strict requirement that particle masses remain constant places strong constraints on multi-fluid formulations. To address this issue, I formulated and implemented a novel algorithm to determine the zero-mass-flux reference frame for each inter-particle flux exchange. This development makes it possible to fully exploit the advantages of the MFM framework, such as accurate shock capturing and reduced numerical diffusion, when modeling dust dynamics.

Future perspectives

The theoretical and numerical advances presented in this thesis open several promising avenues for future work, many of which remain largely unexplored in hydrodynamical simulations. A natural next step is to apply the methods developed here to simulations of astrophysical systems, in order to investigate both the dynamics of dust and its feedback on the evolution of the gas. One class of applications includes isolated or cosmological galaxy simulations, where the dust dynamics module could be coupled with models of dust formation, destruction, and processing in the interstellar medium. This would allow for a more complete understanding of the interplay between dust and gas across galactic scales, including the effects on star formation, chemical enrichment, and feedback-driven outflows. Another important direction is the simulation of molecular cloud collapse, tracking the evolution of dust from galactic scales down to the formation of protoplanetary disks. Differential dust–gas dynamics plays a central role in shaping the early stages of disk formation and the resulting grain-size distribution. While simulations of single protostellar cores (Bate & Lorén-Aguilar, 2016; Lebreuilly et al., 2020) have shown the importance of modeling dust dynamics at these scales rather than assuming tight dust–gas coupling, a gap remains at larger scales. Zoom-in simulations of giant molecular cloud (GMC) collapse by Kuffmeier et al. (2017) highlighted that large-scale heterogeneity affects disk formation, yet did not include dust. More recent frameworks, such as STARFORGE (Grudić et al., 2021), have incorporated dust as Lagrangian particles (Soliman et al., 2024a,b), demonstrating the relevance of modeling both the full grain size distribution and radiation pressure effects on dust. However, these models still lack treatments for dust fragmentation and coagulation, which are challenging to implement in particle-based schemes. The methods developed in this thesis could address this gap, providing a framework to track the full dynamical and

physical evolution of dust from GMCs down to individual protoplanetary disks.
From a modeling perspective, several extensions of the current framework are possible. Incorporating dust fragmentation and coagulation models would allow for a self-consistent evolution of the grain size distribution. These approaches range from high-accuracy Discontinuous – Galerkin methods that evolve the grain size distribution conservatively (Lombart & Laibe, 2021; Lombart et al., 2024), to faster, more versatile methods that are simpler to implement in hydrodynamical simulations and incur a lower computational cost (Stammler & Birnstiel, 2022; Pfeil et al., 2024).
Additionally, alternative dust-dynamics models that explicitly account for particle crossing trajectories could be explored, in order to assess the role of interpenetrating dust flows in regimes characterized by weak dust–gas coupling or high dust-to-gas mass ratios. For example, Vié et al. (2015) developed a realizable Riemann solver for dust dynamics that can, to some extent, mimic particle crossing trajectories within Eulerian fluid simulations. Overall, these developments would broaden the applicability of the methods introduced in this thesis, bridging the gap between fluid-based and particle-based descriptions of dust. They would enable simulations that are both physically more accurate and capable of exploring regimes previously inaccessible to traditional dust–gas modeling approaches.

Bibliography

Andrews, S. M., Huang, J., Pérez, L. M., et al. 2018, ApJ, 869, L41

Asano, R. S., Takeuchi, T. T., Hirashita, H., & Inoue, A. K. 2013, Earth, Planets and Space, 65, 213

Ascher, U. M., Ruuth, S. J., & Spiteri, R. J. 1997, Applied Numerical Mathematics, 25, 151, special Issue on Time Integration

Ballabio, G., Dipierro, G., Veronesi, B., et al. 2018, MNRAS, 477, 2766

Barber, C. B., Dobkin, D. P., & Huhdanpaa, H. 1996, ACM Trans. Math. Softw., 22, 469–483

Bate, M. R. & Lorén-Aguilar, P. 2016, Monthly Notices of the Royal Astronomical Society, 465, 1089

Beck, A. M., Murante, G., Arth, A., et al. 2016, MNRAS, 455, 2110

Benítez-Llambay, P., Krapp, L., & Pessah, M. E. 2019, ApJS, 241, 25

Binkert, F. 2023, MNRAS, 525, 4299

Birnstiel, T. 2024, Annual Review of Astronomy and Astrophysics, 62, 157

Bohren, C. F. & Huffman, D. R. 1998a, Absorption and Scattering by a Sphere (John Wiley Sons, Ltd), 82–129

Bohren, C. F. & Huffman, D. R. 1998b, Particles Small Compared with the Wavelength (John Wiley Sons, Ltd), 130–157

Cacciapuoti, L., Testi, L., Maury, A. J., et al. 2025, A&A, 700, A188

Cartwright, A., Stamatellos, D., & Whitworth, A. P. 2009, MNRAS, 395, 2373

Cheng, N.-S. 2009, Powder Technology, 189, 395

Cherchneff, I. 2014, arXiv e-prints, arXiv:1405.1216

Cullen, L. & Dehnen, W. 2010, MNRAS, 408, 669

Dehnen, W. & Aly, H. 2012, MNRAS, 425, 1068

Di Criscienzo, M., Dell'Agli, F., Ventura, P., et al. 2013, MNRAS, 433, 313

Dolag, K., Vazza, F., Brunetti, G., & Tormen, G. 2005, MNRAS, 364, 753

Draine, B. T., Aniano, G., Krause, O., et al. 2014, ApJ, 780, 172

Draine, B. T. & Li, A. 2001, The Astrophysical Journal, 551, 807

Draine, B. T. & Li, A. 2007, The Astrophysical Journal, 657, 810

Drążkowska, J., Bitsch, B., Lambrechts, M., et al. 2023, in Astronomical Society of the Pacific Conference Series, Vol. 534, Protostars and Planets VII, ed. S. Inutsuka, Y. Aikawa, T. Muto, K. Tomida, & M. Tamura, 717

Du, Q., Faber, V., & Gunzburger, M. 1999, SIAM Review, 41, 637

Dullemond, C. P. & Monnier, J. D. 2010, ARA&A, 48, 205

Epstein, P. S. 1924, Phys. Rev., 23, 710

Ercolano, B. & Pascucci, I. 2017, Royal Society Open Science, 4, 170114

Fitzpatrick, E. L. & Massa, D. 2007, ApJ, 663, 320

Gaburov, E. & Nitadori, K. 2011, MNRAS, 414, 129

Galliano, F., Galametz, M., & Jones, A. P. 2018, ARA&A, 56, 673

Groenewegen, M. A. T. 1997, A&A, 317, 503

Groth, F., Steinwandel, U. P., Valentini, M., & Dolag, K. 2023, MNRAS, 526, 616

Grudić, M. Y., Guszejnov, D., Hopkins, P. F., Offner, S. S. R., & Faucher-Giguère, C.-A. 2021, Monthly Notices of the Royal Astronomical Society, 506, 2199

Hendrix, T. & Keppens, R. 2014, A&A, 562, A114

Hildebrand, R. H. 1983, QJRAS, 24, 267

Hirashita, H. 2013, in Proceedings of The Life Cycle of Dust in the Universe: Observations, ed. C. Kemper, A. Andersen, M. Baes, H. Gomez, & D. Watson, 27

Hirashita, H., Il'in, V. B., Pagani, L., & Lefèvre, C. 2021, MNRAS, 502, 15

Höfner, S. & Olofsson, H. 2018, A&A Rev., 26, 1

Hopkins, P. F. 2015, MNRAS, 450, 53

Hoppe, P. & Zinner, E. 2000, Journal of Geophysical Research: Space Physics, 105, 10371

Hu, C.-Y., Zhukovska, S., Somerville, R. S., & Naab, T. 2019, MNRAS, 487, 3252

Huang, P. & Bai, X.-N. 2022, ApJS, 262, 11

Kataoka, A., Tanaka, H., Okuzumi, S., & Wada, K. 2013, AA, 557, L4

Kirchschlager, F., Schmidt, F. D., Barlow, M. J., et al. 2019, Monthly Notices of the Royal Astronomical Society, 489, 4465

Klahr, H. & Schreiber, A. 2021, ApJ, 911, 9

Kotak, R., Meikle, W. P. S., Farrah, D., et al. 2009, The Astrophysical Journal, 704, 306

Krapp, L. & Benítez-Llambay, P. 2020, Research Notes of the American Astronomical Society, 4, 198

Krapp, L., Garrido-Deutelmoser, J., Benítez-Llambay, P., & Kratter, K. M. 2024a, ApJS, 271, 7

Krapp, L., Garrido-Deutelmoser, J., Benítez-Llambay, P., & Kratter, K. M. 2024b, ApJS, 271, 7

Kuffmeier, M., Haugbølle, T., & Nordlund, Å. 2017, ApJ, 846, 7

Laibe, G. & Price, D. J. 2011, MNRAS, 418, 1491

Laibe, G. & Price, D. J. 2012a, MNRAS, 420, 2345

Laibe, G. & Price, D. J. 2012b, MNRAS, 420, 2365

Laibe, G. & Price, D. J. 2014a, MNRAS, 444, 1940

Laibe, G. & Price, D. J. 2014b, MNRAS, 440, 2136

Laibe, G. & Price, D. J. 2014c, MNRAS, 440, 2147

Lanson, N. & Vila, J.-P. 2008, SIAM Journal on Numerical Analysis, 46, 1935

Lebreuilly, U., Commerçon, B., & Laibe, G. 2019, A&A, 626, A96

Lebreuilly, U., Commerçon, B., & Laibe, G. 2020, AA, 641, A112

Lehmann, A. & Wardle, M. 2018, MNRAS, 476, 3185

Lesur, G., Flock, M., Ercolano, B., et al. 2023, in Astronomical Society of the Pacific Conference Series, Vol. 534, Protostars and Planets VII, ed. S. Inutsuka, Y. Aikawa, T. Muto, K. Tomida, & M. Tamura, 465

Leveque, R. J. 2004, Journal of Hyperbolic Differential Equations, 01, 315

Lombart, M., Bréhier, C.-E., Hutchison, M., & Lee, Y.-N. 2024, Monthly Notices of the Royal Astronomical Society, 533, 4410

Lombart, M. & Laibe, G. 2021, MNRAS, 501, 4298

Lynch, E. M., Lovell, J. B., & Sefilian, A. A. 2024, MNRAS, 529, L147

Marin-Gilabert, T., Valentini, M., Steinwandel, U. P., & Dolag, K. 2022, MNRAS, 517, 5971

Mathis, J. S., Rumpl, W., & Nordsieck, K. H. 1977, ApJ, 217, 425

Mauney, C. M. & Lazzati, D. 2018, Molecular Astrophysics, 12, 1

Maurya, P. K., Yadav, V. S., Mahato, B., et al. 2021, Journal of Computational Physics, 446, 110650

McKinnon, R., Vogelsberger, M., Torrey, P., Marinacci, F., & Kannan, R. 2018a, MNRAS, 478, 2851

McKinnon, R., Vogelsberger, M., Torrey, P., Marinacci, F., & Kannan, R. 2018b, MNRAS, 478, 2851

McNally, C. P., Lyra, W., & Passy, J.-C. 2012, ApJS, 201, 18

Meeus, G., Waters, L. B. F. M., Bouwman, J., et al. 2001, AA, 365, 476

Mignone, A., Flock, M., & Vaidya, B. 2019, ApJS, 244, 38

Minissale, M., Dulieu, F., Cazaux, S., & Hocuk, S. 2016, A&A, 585, A24

Mizerová, H. & Tvrdá, K. 2024, arXiv e-prints, arXiv:2404.16665

Monaghan, J. J. & Kocharyan, A. 1995, Computer Physics Communications, 87, 225

Monaghan, J. J. & Lattanzio, J. C. 1985, A&A, 149, 135

Morris, J. P. 1996, PASA, 13, 97

Nakagawa, Y., Sekiya, M., & Hayashi, C. 1986, Icarus, 67, 375

Nittler, L. R. & Ciesla, F. 2016, ARA&A, 54, 53

Nozawa, T. & Fukugita, M. 2013, ApJ, 770, 27

Pareschi, L. & Russo, G. 2005, Journal of Scientific Computing, 25, 129

Pfeil, T., Birnstiel, T., & Klahr, H. 2024, AA, 691, A45

Planck Collaboration, Adam, R., Ade, P. A. R., et al. 2016, AA, 594, A10

Price, D. J. 2008, Journal of Computational Physics, 227, 10040

Price, D. J. 2012, Journal of Computational Physics, 231, 759

Price, D. J. & Federrath, C. 2010, Monthly Notices of the Royal Astronomical Society, 406, 1659

Price, D. J., Wurster, J., Tricco, T. S., et al. 2018, PASA, 35, e031

Probestein, R. F. & Fassio, F. 1970, AIAA Journal, 8, 772

Rosswog, S. & Price, D. 2007, MNRAS, 379, 915

Sarangi, A. & Cherchneff, I. 2015, AA, 575, A95

Schneider, R. & Maiolino, R. 2024, A&A Rev., 32, 2

Schneider, R., Omukai, K., Inoue, A. K., & Ferrara, A. 2006, Monthly Notices of the Royal Astronomical Society, 369, 1437

Sedov, L. I. 1946, Journal of Applied Mathematics and Mechanics, 10, 241

Sedov, L. I. 1959, Similarity and Dimensional Methods in Mechanics

Slavin, J. D., Dwek, E., Mac Low, M.-M., & Hill, A. S. 2020, The Astrophysical Journal, 902, 135

Sod, G. A. 1978, Journal of Computational Physics, 27, 1

Soliman, N. H. & Hopkins, P. F. 2023, MNRAS, 525, 2668

Soliman, N. H., Hopkins, P. F., & Grudić, M. Y. 2024a, ApJ, 974, 136

Soliman, N. H., Hopkins, P. F., & Grudić, M. Y. 2024b, ApJ, 975, 284

Springel, V. 2005, MNRAS, 364, 1105

Springel, V. & Hernquist, L. 2002, MNRAS, 333, 649

Stammler, S. M. & Birnstiel, T. 2022, ApJ, 935, 35

Steinwandel, U. P., Moster, B. P., Naab, T., Hu, C.-Y., & Walch, S. 2020, MNRAS, 495, 1035

Tedeschi-Prades, G., Birnstiel, T., Dolag, K., Ercolano, B., & Hutchison, M. 2025a, The Astrophysical Journal, 994, 267

Tedeschi-Prades, G., Birnstiel, T., Dolag, K., Ercolano, B., & Hutchison, M. 2025b, A&A, 704, A92

Toro, E. 2009, Riemann Solvers and Numerical Methods for Fluid Dynamics: A Practical Introduction

Trumpler, R. J. 1930, PASP, 42, 214

Valiante, R., Schneider, R., Bianchi, S., & Andersen, A. C. 2009, Monthly Notices of the Royal Astronomical Society, 397, 1661

van der Marel, N. 2023, European Physical Journal Plus, 138, 225

Van Leer, B. 1977, Journal of Computational Physics, 23, 276

Vericel, A., Gonzalez, J.-F., Price, D. J., Laibe, G., & Pinte, C. 2021, MNRAS, 507, 2318

Vié, A., Doisneau, F., & Massot, M. 2015, Communications in Computational Physics, 17, 1

Vogelsberger, M., McKinnon, R., O'Neil, S., et al. 2019, MNRAS, 487, 4870

Weingartner, J. C. & Draine, B. T. 2001, ApJ, 548, 296

Wendland, H. 1995, Advances in Computational Mathematics, 4, 389

Yan, H., Lazarian, A., & Draine, B. T. 2004, The Astrophysical Journal, 616, 895

Yang, C.-C. & Johansen, A. 2016, ApJS, 224, 39

Youdin, A. N. & Goodman, J. 2005, ApJ, 620, 459

Zhang, R., Yuan, H., & Chen, B. 2023, ApJS, 269, 6

Zhukovska, Gail, H.-P., & Trieloff, M. 2008, AA, 479, 453

Acknowledgements

I would first like to thank the people who made this work possible: Professors Til Birnstiel, Klaus Dolag, and Barbara Ercolano. Thank you for your continuous support, your invaluable advice, and for believing in this project.
I would also like to extend my thanks to the people with whom I shared these years at the Observatory, and in particular to the members, past and present, of Til's group: Thomas, Tommy, Luca D., Fabian, Andrés, Joanna, Sebastian, Asmita, Anna, and Alex.
Thank you to Geray for making `OpenGadget3` a better place, and to Freddy for patiently answering my countless questions about its inner workings.

Ringrazio anche i miei genitori, che mi hanno sempre sostenuto. Grazie per esservi sempre preoccupati per me, e per essere stati un esempio per tutti questi anni.
Grazie a Santiago e Roberto, per ricordarmi da dove vengo, anche dopo tanti anni lontano da casa.
Grazie a Michele e Alfonso, per esserci sempre quando serve, per una risata in più.
Sopra a tutti, grazie a Giulia. Questo percorso non sarbbe stato bello la metà, senza la tua presenza. Condividerlo è stata la parte più preziosa del percorso.