Numerical Methods for Simulations of Planetary Impacts

THOMAS DANIEL SANDNES

SANDNES, THOMAS DANIEL (2024) Numerical Methods for Simulations of Planetary Impacts. Doctoral thesis, Durham University.

Numerical Methods for Simulations of Planetary Impacts

Thomas Daniel Sandnes

A Thesis presented for the degree of
Doctor of Philosophy

Institute for Computational Cosmology
Department of Physics
Durham University
United Kingdom
November 2024

Abstract

Computational simulations are an invaluable tool for studying the complex processes that shape planetary systems. This includes the impact events, both small and giant, whose lasting effects can be observed throughout the solar system to this day. This work details the development and validation of numerical methods for smoothed particle hydrodynamics (SPH) simulations, and the application of these methods to simulations of planetary impacts.

With traditional SPH formulations, fluid mixing and instability growth are artificially suppressed by spurious surface tension-like effects. The suppression of mixing between dissimilar, stiff materials, such as those used in planetary impact simulations, is especially strong and difficult to alleviate. While various approaches have been developed to mitigate this issue, they often introduce dependencies on specific material properties, rely on targeted or contrived corrections, or require bespoke particle configurations that cannot be maintained throughout the course of typical science simulations. In this thesis, a new SPH formalism is developed. By directly targeting sources of error, this generalised, material-independent approach improves the treatment both of mixing within a single material, for example in an ideal gas, and at interfaces between dissimilar materials. This new SPH scheme is validated in a range of hydrodynamic test simulations, including both standard test cases and planetary-specific scenarios.

These methods are then applied in simulations of planetary giant impacts onto Jupiter to investigate whether this mechanism could be responsible for the planet's observed dilute core. With these new methods, there is a significant increase in the amount of mixing between core and envelope material during impact simulations. Although core material is temporarily diluted, heavy elements settle under gravity and no dilute core is produced in any of the giant impact simulations carried out with different impact speeds, impact angles, internal structures, equations of state, or numerical resolutions.

The capabilities of the newly developed methods are extended to include solid-body mechanics. By incorporating physical models for elasticity, plastic deformation, and fracturing, impacts at much smaller scales can be simulated. The implementation of these models is validated in a range of test simulations.

Declaration

The work in this thesis is based on research carried out at the Department of Physics, Durham University, United Kingdom. No part of this thesis has been submitted elsewhere for any other degree or qualification and it is all my own work unless referenced to the contrary in the text.

Acknowledgements

I am incredibly grateful to my supervisors, Vince Eke and Richard Massey, for all your guidance and support, especially in the tough times when it felt like steps forward and steps back had found a dynamic equilibrium. A special thank you to Jacob Kegerreis for going above and beyond with all your help and enthusiasm at every stage of this work. To Luis Teodoro, thank you for all your much-needed positivity, and to Sergio Ruiz-Bonilla, thank you for setting me on the path that ultimately led to the majority of the pages in this thesis. Thank you all for helping me finding the method(s) in the madness.

Thank you to Matthieu Schaller and the entire SWIFT team for all the work you have done, and continue to do, that has made everything in this thesis both possible and enjoyable.

To Ben D., Ben M., Ben R., Jack, Luca, Matt, Morgan, and Niall – I am incredibly fortunate to have friends like you. And to Dan, Jon, and Simon, who helped me stay on track through the trilogy of tricky problems.

Most of all, thank you to my family and Rachel. Without your love and support none of this would have been possible. I am more grateful than you could ever know.

This thesis was supported by STFC grants ST/T506047/1 and ST/V506643/1.

Contents

List of Figures

List of Tables

For Rachel

CHAPTER 1

Introduction

Planets are born within disks of dust and gas. Micron-sized dust particles grow by coagulation to form larger objects (Safronov, 1972). Over time, planetesimals of order ∼10 km are thought to form from collisions (Benz, 2000) and gravitational collapse (Goldreich & Ward, 1973; Youdin & Shu, 2002). Ultimately, planets form through the successive impacts between planetesimals, and the accretion of gas from the surrounding disk (Armitage, 2020). After a few million years, the remaining gas and dust in the protoplanetary disk dissipate primarily through photoevaporation due to radiation from the central star (Haisch Jr et al., 2001), leaving behind a planetary system that more closely resembles the Solar System as we know it today. While significant progress has been made in studying the mechanisms of planet formation, these processes are complex and therefore, in many instances, are still not well understood. Although our Solar System is no longer as tumultuous as it once was, studying the present-day system offers insights into its evolutionary history and the processes that shaped it. Additionally, with the advancement of observational capabilities, we can now glimpse into the environments in which planets form, and observe these processes directly in dusty disks around distant, young stars (ALMA Partnership et al., 2015).

Impact events have left clues to the conditions and dynamics of the evolving Solar System, shaping the surfaces of planets and, in some cases, fundamentally altering their evolutionary trajectory. These impacts vary widely in scales of size and energy, from those in which asteroids peppered the Moon's surface with craters (Baldwin, 1949; Gilbert, 1893), to catastrophic giant impacts, like the one thought to have formed the Moon itself (Cameron & Ward, 1976; Hartmann & Davis, 1975).

Since the mechanisms that govern the formation of planets are highly coupled and nonlinear, it is often most effective to employ computational simulations to investigate these processes (Fermi et al., 1955). Rapid advancements in technology and numerical methods have driven a substantial progression in the capabilities of computational simulations for planetary impact applications, from the pioneering smoothed particle hydrodynamics simulations of Benz et al. (1986) to present-day simulations with up to $\sim 10^{5.5}$ times higher resolution, such as those presented in this thesis.

As the field of planetary science rapidly expands beyond our own Solar System, with observations of planetary systems around distant stars, we are in an exciting era for studying the impacts that shape planetary systems (Kenworthy et al., 2023; Mayor & Queloz, 1995). The work of this thesis aims to significantly advance the numerical methods used for planetary impact simulations, additionally incorporating physical models that allow simulations across a range of scales, spanning impacts between giant planets to small cratering simulations where solid-body strength physics can dominate. Understanding sensitivities in simulation outcomes to numerical aspects of the methods employed is crucial for the interpretation of results. This focus on furthering understanding of, and addressing shortcomings in, simulation methodologies underpins the work in this thesis, and motivates the presented improvements to numerical models that aim to lead to more physically accurate planetary impact simulations.

1.1 Numerical methods for simulating fluids

Computational simulations are an invaluable tool for studying the inherently complex behaviour of fluids. A variety of methods have been developed across science and engineering disciplines to model fluid dynamics, each with strengths that depend on the application and the properties of the physical system being modelled. These methods are broadly classified into two categories: Eulerian methods, where the simulated fluid moves through a stationary grid; and Lagrangian methods, where interpolation points move with the fluid, tracing its streamlines. Additionally, there are methods that bridge the gap between these approaches, such as methods that use a moving mesh (Springel, 2010b).

Within these categories, further distinctions can be made based on how the Navier–Stokes equations (or Euler equations, for inviscid flow) are solved. For instance, Eulerian methods can be subdivided into:

- finite difference methods (FDM): solutions to partial differential equations are approximated at grid-point positions by expressing derivatives as differences of quantities over discrete intervals of space and time (LeVeque, 2007);
- finite volume methods (FVM): fluxes of fluid quantities through surfaces separating volumes associated with grid-points are used to evolve the fluid (Moukalled et al., 2016);
- finite element methods (FEM): the simulation domain is discretised into a set of elements (e.g. triangles in 2D or tetrahedra in 3D). Within each element, fluid properties are approximated by linear combinations of basis functions, which are calculated from values sampled at element vertices (Zienkiewicz & Taylor, 2005).

As for Lagrangian methods, smoothed particle hydrodynamics (SPH) is the most widely used approach and is central to the work of this thesis. Other Lagrangian methods are often more specialised, such as the Particle-In-Cell (PIC) method used for plasma physics (Tskhakaya et al., 2007) and vortex methods (Mimeau & Mortazavi, 2021) for vorticity-dominated incompressible flows.

1.2 Smoothed particle hydrodynamics

Smoothed particle hydrodynamics, first developed by Lucy (1977) and Gingold & Monaghan (1977), is frequently utilised across a range of applications spanning astrophysics (Benz, 1988; Monaghan, 1992; Springel, 2010a) and engineering (Libersky et al., 1993; Lind et al., 2020). In astrophysics, it is used in particular for its geometry-independent adaptive resolution, inherent conservation properties, and elegant coupling with gravity solvers (Springel, 2005). For engineering applications, it offers advantages in the treatment of dynamic free surfaces, fluid–structure interactions, and in simulating multiphase flow (Violeau & Rogers, 2016; Wang et al., 2016). This is in addition to its relatively simple construction, numerical stability, and low computational cost. SPH is the most wide-spread method for simulating impacts between planetary bodies, since all the benefits outlined above apply in these scenarios. In this thesis, my motivation and methods focus on conservative, fully-compressible, gravity-coupled SPH schemes, where particles have unchanging masses and material types throughout a simulation.

Two key concepts characterise SPH: the representation of a fluid as a discrete set of interpolation points, or ‘particles’, that move with the fluid velocity; and the use of a kernel function to estimate fluid fields and their gradients at particle positions, by interpolation over neighbouring particles (Price, 2012). However, specific errors are introduced with the assumptions that underpin these core concepts. The discretisation of the continuous underlying fluid results in leading-order error in the momentum equation, which is sensitive to disorder in the local particle distribution (Read et al., 2010). Additionally, the use of an extended kernel in the traditional, integral form of the SPH density estimate leads to inadvertent smoothing of interpolated densities (Monaghan, 1985; Price, 2012). In regions where variations in the underlying density field are not well resolved by the instantaneous particle configuration, this can lead to the calculation of spurious particle pressures, and subsequently to spurious pressure gradients that are used in the equations of motion.

These errors combine particularly strongly at density discontinuities in simulations where a fluid is represented by particles of fixed, equal mass – which is the norm for many science applications. In such a case, a density discontinuity constitutes a sharp change in particle spacing. Both discretisation and kernel smoothing error combine to give rise to a spurious surface tension-like effect that greatly suppresses both the mixing of fluid across the interface and the growth of instabilities that would act to drive turbulent mixing (Agertz et al., 2007). This is a well-established shortcoming of SPH. In this thesis I aim to construct an SPH formulation that directly addresses these sources of error, thereby improving the treatment of mixing in simulations.

As well as its use as a tool to simulate fluids, the SPH framework can be adapted to model material strength physics to simulate solid-state mechanics (Libersky & Petschek, 1991). An elastic–perfectly-plastic constitutive model is often used, in which the accumulation of stress is described by Hooke's law for a continuous medium and is limited by a yield stress criterion to capture plastic deformation. These models can be combined with methods to simulate the growth of fractures by the accumulation of damage (Benz & Asphaug, 1994, 1995; Jutzi, 2015) and models for porosity (Jutzi et al., 2008). These methods have been used in the context of planetary impacts to simulate impacts across a wide range of scales, including impacts onto asteroids in the context of the DART mission (Raducan & Jutzi, 2022), the tidal disruption of the comet Shoemaker–Levy 9 (Asphaug & Benz, 1996), and giant impacts at the scale of Mars (Ballantyne et al., 2023; Emsenhuber et al., 2018).

In the following sections, I describe the key constituent components of SPH. We use an SPH formulation based on that of Price (2012), summarised in Appendix C, as a basis for our discussion and for comparisons throughout. Additionally, I take this as an opportunity to introduce the nomenclature and notation that is used throughout this thesis. An additional glossary of notation is included in Appendix A. The strength models considered in this thesis are introduced in Chapter 4, along with results from relevant test simulations.

1.2.1 Kernel interpolation and the SPH density estimate

Kernel interpolation theory forms the framework for SPH estimates of fluid fields and their gradients. In particular, the integral form of the density estimate is a core component of many SPH schemes (Price, 2012), by which a smoothed density field at the position of particle i, $\langle\rho_i\rangle$, can be reconstructed from the local spatial distribution of neighbouring particles j, their masses, m_j, and a kernel function, W_{ij} (described below), via

$$\langle\rho_i\rangle = \sum_j m_j W_{ij} \; . \tag{1.1}$$

Throughout the governing equations of these SPH schemes, the interpolated density, $\langle\rho_i\rangle$, is used as an estimate of the underlying density field at the positions of particles, $\rho(\mathbf{r}_i)$. This density estimate is a specific application of kernel interpolation, which in general can be used to reconstruct an arbitrary field, F, from its value sampled at the positions of particle neighbours, via

$$\langle F_i\rangle = \sum_j F_j W_{ij} V_j \; , \tag{1.2}$$

where V_j are volume elements of particle j. In Eqn. 1.1, volume elements are taken to be $V_j = m_j/\rho_j$. Kernel interpolation can also be used in estimates of the gradient[1] of F,

$$\left\langle \left.\frac{dF}{d\mathbf{r}}\right|_i \right\rangle = \sum_j F_j \left.\frac{dW}{d\mathbf{r}}\right|_{ij} V_j \; , \tag{1.3}$$

such as in the calculation of pressure gradients and velocity divergences for the SPH equations of motion.

The smoothing kernel, $W(\mathbf{r}-\mathbf{r}', h(\mathbf{r}))$, is a weighting function with radial extent characterised by the smoothing length h. $W(\mathbf{r}-\mathbf{r}', h(\mathbf{r}))$ approaches a delta function in the limit $h(\mathbf{r}) \to 0$. Traditionally, W is a positive function with approximately a truncated Gaussian-like shape; the kernel is typically normalised, and spherically symmetric, as this ensures the exact interpolation of linear fields in the continuum limit of kernel sampling (number of neighbours, $N \to \infty$). For a particle pair i, j: $W_{ij} \equiv W(\mathbf{r}_{ij}, h_i) \equiv W(\mathbf{r}_i - \mathbf{r}_j, h(\mathbf{r}_i))$, where $\mathbf{r}_{ij} \equiv \mathbf{r}_i - \mathbf{r}_j$. Subscripts denote

---

[1] We make the choice of notation, here and throughout, to express kernel gradients as total derivatives rather than with "∇" which is often used to imply derivatives with fixed smoothing length. In later sections, this allows us to more easily distinguish between gradient estimates with and without grad-h terms (Hopkins, 2013).

quantities either sampled at the position of, or associated with, a particle. Kernels with a compact support, $H \equiv H(h)$ such that $W(r > H) = 0$, are used to limit the number of neighbours to a finite number. We adopt the convention of defining the smoothing length h as twice the standard deviation of the kernel[2] (Dehnen & Aly, 2012). This relates the smoothing length to the compact support by a constant multiplication factor H/h.

In the SPH construction presented here, a particle's smoothing length is evaluated iteratively to satisfy

$$h_i = \eta_{\rm kernel} \left(\frac{1}{\sum_j W(\mathbf{r}_{ij}, h_i)} \right)^{1/d} , \tag{1.4}$$

where d is the spatial dimensionality of the simulation and $\eta_{\rm kernel}$ is a chosen constant. Eqn. 1.4 ensures that particles across the simulation have an approximately constant number of neighbours, determined by the form of the kernel function and the choice of $\eta_{\rm kernel}$.

We use the Wendland C^2 kernel (Wendland, 1995) for the primary simulations of this thesis:

$$W_{\rm WC2}\left(|\mathbf{r} - \mathbf{r}'|, H(h)\right) = \begin{cases} \dfrac{C}{H^d} \left(1 - \dfrac{|\mathbf{r} - \mathbf{r}'|}{H}\right)^4 \left(1 + 4\dfrac{|\mathbf{r} - \mathbf{r}'|}{H}\right) & \text{for } |\mathbf{r} - \mathbf{r}'| < H \, , \\ 0 & \text{otherwise} \, , \end{cases} \tag{1.5}$$

with $\eta_{\rm kernel} = 1.487$ ($\sim$100 neighbours for $d = 3$) (Dehnen & Aly, 2012). Here $C = 21/(2\pi)$ is the normalisation constant for the Wendland C^2 kernel in 3D. Higher-order kernels can reduce error but require greater numbers of neighbours, which can come at a significant cost to code speed. This kernel offers a suitable compromise between improved accuracy and fast simulation run-time, a relevant consideration for science applications. In Appendix D I demonstrate the effect of

---

[2] Despite the ubiquitous use of the nomenclature and notation of the "smoothing length, h", different definitions are frequently used for both the relationship between H and h, and the method used to calculate h (in our case, Eqn. 1.4). Although the differences are subtle, we draw attention to this as an example of the difficulty of one-to-one comparisons between simulation codes, especially as methods become increasingly complex.

the choice of kernel function in hydrodynamic test simulations.

1.2.2 SPH equations of motion

The equations of motion govern the kinematic and thermodynamic evolution of SPH particles. The Euler equations are used as the basis for the SPH equations of motion for inviscid fluids. These consist of the continuity equation, momentum equation, and energy equation, which are closed by an "equation of state". The general, thermodynamically consistent SPH equations of motion, where we additionally use the same kernel function across the equations, take the form (Hopkins, 2013; Read et al., 2010)

$$\frac{d\rho_i}{dt} = \sum_j m_j \frac{\zeta_i}{\zeta_j} \mathbf{v}_{ij} \cdot \left.\frac{dW}{d\mathbf{r}}\right|_{ij}, \tag{1.6}$$

$$\frac{d\mathbf{v}_i}{dt} = -\sum_j m_j \left(\frac{P_i}{\rho_i^2} \frac{\xi_i}{\xi_j} + \frac{P_j}{\rho_j^2} \frac{\xi_j}{\xi_i} \right) \left.\frac{dW}{d\mathbf{r}}\right|_{ij}, \tag{1.7}$$

$$\frac{du_i}{dt} = \frac{P_i}{\rho_i^2} \sum_j m_j \frac{\zeta_i}{\zeta_j} \mathbf{v}_{ij} \cdot \left.\frac{dW}{d\mathbf{r}}\right|_{ij}, \tag{1.8}$$

where particle densities, ρ_i, velocities, $\mathbf{v}_i$, and specific internal energies, u_i, are evolved in time based on gradients of pressure, P_i, and velocity divergences calculated using the relative velocity of particle pairs $\mathbf{v}_{ij} \equiv \mathbf{v}_i - \mathbf{v}_j$. The free functions ζ and ξ are introduced in the process of discretisation. An SPH scheme that explicitly conserves energy and momentum requires antisymmetric kernel gradient terms in the exchange of particle pairs i and j. The integral form of the density estimate (Eqn. 1.1) is equivalent to the differential form (Eqn. 1.6) in the continuum limit, for $\zeta_i = \zeta_j$ (Read et al., 2010).

1.3 Sources of error in SPH

In this section I describe key sources of error in the SPH formulation, whose reduction is central to this thesis. I summarise approaches that have been taken previously to address these, and their limitations.

1.3.1 Kernel smoothing error

A fluid field reconstructed using an extended kernel with $h \neq 0$ will be affected by smoothing error, even when sampled in the continuum limit (Price, 2012). In the continuum limit, a reconstructed field, $\langle F \rangle$, is the convolution of the underlying field, F, with a smoothing kernel W,

$$\langle F(\mathbf{r}, h) \rangle = \int F(\mathbf{r}')\, W(\mathbf{r} - \mathbf{r}', h)\, dV' \,. \tag{1.9}$$

Eqn. 1.2 is the discretised form of this equation. Assuming a continuous, infinitely differentiable field F, we can Taylor expand about the point $\mathbf{r}$ to give

$$\begin{aligned} \langle F(\mathbf{r}, h) \rangle \;=\; & F(\mathbf{r}) \cancelto{1}{\int W(\mathbf{r} - \mathbf{r}', h)\, dV'} \;+\; \left.\frac{dF}{dr^{\alpha}}\right|_{\mathbf{r}} \cancelto{0}{\int (\mathbf{r}' - \mathbf{r})^{\alpha}\, W(\mathbf{r} - \mathbf{r}', h)\, dV'} \\ & + \; \frac{1}{2} \left.\frac{d^2F}{dr^{\alpha} dr^{\beta}}\right|_{\mathbf{r}} \int (\mathbf{r}' - \mathbf{r})^{\alpha}\, (\mathbf{r}' - \mathbf{r})^{\beta}\, W(\mathbf{r} - \mathbf{r}', h)\, dV' \;+\; \ldots \end{aligned} \tag{1.10}$$

$$= F(\mathbf{r}) \;+\; \frac{1}{2} \left.\frac{d^2F}{dr^{\alpha} dr^{\beta}}\right|_{\mathbf{r}} \int (\mathbf{r}' - \mathbf{r})^{\alpha}\, (\mathbf{r}' - \mathbf{r})^{\beta}\, W(\mathbf{r} - \mathbf{r}', h)\, dV' \;+\; \ldots \,, \tag{1.11}$$

where Greek letter superscripts correspond to spatial dimensions, and like indices are summed over (Price, 2012; Sigalotti et al., 2019). We separate the first two terms in Eqn. 1.10 to demonstrate that, in the continuum limit, the choice of a normalised, spherically symmetric kernel results in the zeroth- and first-order integrals of the expansion taking values 1 and $\mathbf{0}$ respectively.

In the continuum limit, kernel interpolation will only reproduce $F(\mathbf{r})$ without error if the integrals of the second and higher order terms are all equal to zero. Due to the assumed symmetry properties of the kernel, integrals in odd terms of the expansion are trivially equal to 0, while in general even terms will be non-zero. For a positive kernel these non-zero terms act to smooth the reconstructed field. Although the integrals in Eqn. 1.11 will be of order h^2 and higher powers of h, with exponents corresponding to the term of the expansion (Monaghan, 1985), the errors become significant in regions where second and higher order derivatives of the underlying field are large over length scales of h (Violeau & Fonty, 2019). This is, in particular, the case for an underlying field that approaches a discontinuity relative to h-length scales. A discontinuity, where the field is not differentiable, will inevitably be erroneously smoothed by kernel interpolation.

The integral SPH density estimate, Eqn. 1.1, is an example of the discrete form of Eqn. 1.9. Through Eqn. 1.11, we see how a quantity calculated by kernel interpolation in this way will experience smoothing error when the underlying field varies sharply over h-length scales. At density discontinuities, smoothing of the density field, while retaining a sharp jump in specific internal energy, leads to spurious pressures that contribute to surface tension-like effects that impede particle mixing across the interface.

Many past approaches to address kernel smoothing error at contact discontinuities explicitly assume the use of a single, ideal gas EoS. In particular, we note the use of artificial conduction for this purpose, by which particle internal energies are smoothed over a similar length scale to the inadvertent density smoothing (Price, 2008). This requires thermodynamic behaviour such that smooth density and internal energy fields result in a smooth pressure field. Therefore, this cannot reliably improve the treatment of interfaces between dissimilar materials, represented by different EoS. Alternatively, methods that use modified density estimates, weighted by a simple thermodynamic quantity such as specific internal energy, also assume a simple relationship between density and internal energy at constant pressure, typically the inversely proportional relationship of an ideal gas (Read et al., 2010; Ritchie & Thomas, 2001).

Dealing with kernel smoothing errors at interfaces between arbitrarily different materials is more challenging since the simplicity of the ideal gas equation cannot be exploited. A boundary between dissimilar materials in thermal and pressure equilibrium will in general result in a density discontinuity, so these problematic scenarios occur frequently in simulations with multiple materials. Additionally, the surface tension-like effects caused by the density smoothing are particularly strong for “stiff” EoS, for which small changes in particle densities can result in large changes in calculated pressures (Melosh, 1989).

Methods to improve the treatment of material interfaces in fully-compressible SPH formulations have been explored in the context of planetary impacts (Benz et al., 1989), where density discontinuities between multiple, stiff materials are common and can evolve across a range of thermodynamic phase space throughout the course of a single simulation. The treatment of discontinuous free surface interfaces is also important in this context (Reinhardt & Stadel, 2017). Hosono et al. (2013) present a “density independent SPH” (DISPH (Saitoh & Makino, 2013)) scheme adapted for use with multiple materials. Here, rather than being calculated from particle masses and densities, volume elements are based on functions of pressure that are evolved in time and recalculated to satisfy kernel normalisation in an additional iterative step (Hosono et al., 2016). Although this approach prevents spurious pressures at density discontinuities, specifically in regions with otherwise continuous pressures, the extension of this method to arbitrary EoS leads to material-dependent volume elements that intricately depend on fluid thermodynamics. Pearl et al. (2022) present an advanced scheme that, among other improvements, makes use of Riemann solvers (Inutsuka, 2002) and an optional slip condition at material interfaces. Their choice of material-dependent density estimate effectively smooths volume rather than density at material interfaces, in simulations where particles of the same material have equal masses. If particles are deliberately set up with equal volumes, then this density estimate will significantly reduce both kernel-smoothing and discretisation errors. This improvement is evident in Pearl et al. (2022)’s mixing tests, where particles start on a single, ordered grid. However, such a bespoke particle configuration is not representative of many science applications. Therefore, these tests do not validate mixing with their methods for general cases with emergent and evolving density discontinuities, where particle configurations cannot easily be controlled in this way throughout the simulation.

This approach of addressing smoothing error through the choice of particle masses is also taken by Deng et al. (2019), who demonstrate enhanced mixing in their meshless finite-mass (MFM) (Gaburov & Nitadori, 2011; Hopkins, 2015; Ivanova et al., 2013; Vila, 1999) simulations of planetary giant impacts. Although

MFM includes Riemann solvers and more advanced gradient estimates that can improve on standard SPH formulations, densities are still calculated with an interpolated estimate that in this case smooths volume, and therefore is still subject to kernel smoothing error. Additionally, a range of SPH modifications specific to material boundaries, rather than arbitrary density discontinuities, have also been developed (Reinhardt et al., 2020; Ruiz-Bonilla et al., 2022; Woolfson, 2007). All the methods discussed here that address density smoothing directly, rather than in the construction of initial conditions, rely on EoS- or material-dependent treatments in, for example, the calculation of volume elements and density estimates.

1.3.2 Discretisation error

The kernel smoothing errors discussed above are in addition to, and separate from, errors introduced by discretisation (Price, 2012; Spreng et al., 2020). Discretisation errors manifest themselves both through the choice of free functions in the equations of motion – affecting how closely Eqns. 1.6–1.8 approximate their continuous Euler equation equivalents – and through the imperfect sampling of the kernel by a finite number of particle neighbours, i.e. in the discretisation of integrals like Eqn. 1.9.

The use of a normalised, spherically symmetric kernel leads to the exact reconstruction of linear fields in the continuum limit by Eqn. 1.9, as the higher-order derivatives in Eqn. 1.11 are zero by construction. However, in the process of discretisation of the fluid into a finite set of particles, the conditions

$$\sum_j W_{ij} V_j = 1 \ , \tag{1.12}$$

$$\sum_j \mathbf{r}_{ij} W_{ij} V_j = \mathbf{0} \ , \tag{1.13}$$

are no longer enforced. The exact reconstruction of fluid fields is therefore lost, even to zeroth order. The amount of discretisation error is a function of the disorder in the local particle distribution. This also applies to gradient estimates, such as those used in the equations of motion. Furthermore, in the equations of motion, gradient estimates are typically modified to enforce conservation, so generally deviate further from exact reproduction of underlying linear fields.

In SPH simulations where a fluid is represented by particles of equal mass, a density discontinuity constitutes a sharp change in particle spacing and thus large local anisotropies in particle distribution. This leads to discretisation error also playing a considerable role in suppressing mixing at density discontinuities (Read et al., 2010).

Using higher-order kernel functions with more particle neighbours will generally reduce error (Dehnen & Aly, 2012), and choices of free functions in the generalised form of the equations of motion can be exploited to mitigate zeroth-order error (Read et al., 2010; Wadsley et al., 2017). In conjunction with these, improved gradient estimates from, for example, reproducing kernels (Frontiere et al., 2017; Liu & Jun, 1998; Liu et al., 1995) or integral-based gradient estimates (García-Senz et al., 2012; Rosswog, 2015, 2020) have been demonstrated to improve the treatment of fluid mixing and instability growth. These methods have no dependence on material or equation of state (EoS) in their construction or underlying assumptions. We therefore make use of some of these methods in the work of this thesis.

1.4 Thesis motivations and outline

The first goal of this thesis is to develop a new SPH formulation to improve the treatment of mixing in our simulations, in particular for the complex equations of state used to simulate planetary impacts. The constitutive equations of this formulation and validation test simulations are presented in Chapter 2. We approach this problem by directly targeting sources of kernel smoothing error and discretisation error, resulting in a generalised, material-independent formulation that improves the treatment both of discontinuities within a single material, for example in an ideal gas, and of interfaces between dissimilar materials. In contrast with contemporary methods, we achieve these improvements while maintaining sharp discontinuities; without introducing additional equation of state dependence in, for example, particle volume elements; without contrived or targeted corrections. Additionally, we focus on the more challenging cases with particles of equal mass across the simulation to validate the methods for typical simulations of planetary impacts, as well as use-cases relevant across applications spanning astrophysics and engineering, where particles are free to evolve over a large range of density scales, or where emergent and evolving density discontinuities cannot easily be corrected by choosing bespoke particle masses in the initial conditions.

Next, I apply this SPH scheme to a problem that tests its capabilities for simulations of planetary impacts. Measurements of Jupiter's gravitational moments by the Juno spacecraft have lead to models of the planet's interior that suggest that there is a smooth compositional gradient between the planet's core of heavy elements and its primarily hydrogen–helium envelope, rather than a sharp boundary (Wahl et al., 2017). The disruption of a differentiated core by an extreme, head-on giant impact has been suggested as a possible formation mechanism for the planet's dilute core (Liu et al., 2019). In Chapter 3 I present a study of simulations of giant impacts onto Jupiter using our newly developed hydrodynamical methods to investigate whether a giant impact could have mixed the planet's core into a stable, diluted configuration.

Finally, I extend the capabilities of our numerical methods to a regime of smaller scales, where the representation of planetary materials as fluids breaks down. I implement material strength physics into our newly developed SPH formulation. This includes models for elasticity, plasticity, and fracturing. In Chapter 4, I present results from test simulations to validate our methods with these physical models.

CHAPTER 2

REMIX SPH

Here, we present the REMIX (Reduced Error MIXing) SPH scheme. REMIX is constructed with the following goals in mind: (1) to improve the treatment of density discontinuities and mixing in simulations with both one and multiple EoS, by directly addressing sources of error in traditional SPH methods; (2) to be able to achieve this for simulations with particles of equal mass, as well as less challenging configurations; (3) to retain the key characteristics of the SPH formalism; (4) to introduce no additional EoS dependence in, for example, volume elements or density estimates; (5) for computational efficiency, to require no more than three loops over particle neighbours and no additional iterative steps compared with the traditional formulation. An implementation of REMIX is publicly available as part of the open-source SWIFT code[1] (Schaller et al., 2024).

This Chapter is structured as follows: in §2.1, we describe the methods used in practice to run our simulations; in §2.2, we present each component of the REMIX SPH scheme; in §2.3, we validate REMIX in a range of hydrodynamic test simulations; and we summarise our findings in §2.4.

---

[1] SWIFT is in open development including extensive documentation and examples at swiftsim.com.

2.1 Methods

2.1.1 Equations of state

The equation of state (EoS) characterises the thermodynamic behaviour of a material. In SPH simulations, hydrodynamical evolution is tied directly to pressures and sound speeds, calculated through the EoS. Many applications in astrophysics use simulations with only a single, ideal gas EoS. However, in some cases, multiple EoS are required to simulate dissimilar materials or phases, such as for planetary impacts, where EoS are often highly complex (Melosh, 1989). The improvements offered by the REMIX SPH scheme are EoS-independent, and so our methods can be applied effectively to these simulations, as well as other applications with multiphase fluids.

For the hydrodynamic test simulations presented in §2.3, we validate the REMIX scheme using both ideal gases and more complex EoS. For ideal gas simulations, the adiabatic index, γ, is problem-specific and chosen to draw comparisons with past work. For simulations using more complex materials, we use EoS typically used for planetary impact simulations. In most of these tests, we consider iron and rock in conditions representative of the core–mantle boundary in an Earth-like planet. We use the updated ANEOS $Fe_{85}Si_{15}$ and forsterite EoS for these materials, respectively (Stewart et al., 2020). For simplicity, we hereafter refer to these as "iron" and "rock". In §2.3.9, we also consider a Jupiter-like planet. For these simulations, we use the hydrogen–helium EoS from Chabrier & Debras (2021), with a helium mass fraction of $Y = 0.245$, and the AQUA EoS from Haldemann et al. (2020) to represent heavy elements or ice.

We note that in the simulations we present in this chapter, these materials are treated as fluids without physical viscosities or strength properties.

2.1.2 The SWIFT code

SWIFT is a state-of-the-art, open-source hydrodynamics and gravity code that specialises in SPH simulations for planetary applications as well as galaxy formation and cosmology (Kegerreis et al., 2019; Schaller et al., 2024). By using task-based parallelism, asynchronous communications, and graph-based decomposition of the work between compute nodes, SWIFT can perform high-resolution simulations efficiently on modern high-performance computing architectures (Schaller et al., 2016). REMIX is fully integrated into and was developed using the SWIFT code, and is therefore publicly available[2]. All simulations presented here were carried out using the SWIFT code. Algorithms used for the neighbour-finding, time-stepping, and gravity in our simulations are detailed in Schaller et al. (2024) and are used identically for simulations with both REMIX and traditional SPH.

2.2 The REMIX formulation

In this section, we detail the constitutive equations of the REMIX SPH scheme[3]. We improve the treatment of mixing by directly addressing the sources of SPH error discussed in §1.3. By targeting both smoothing and discretisation error, we alleviate spurious surface tension-like effects at density discontinuities, including in the more challenging cases with equal-mass particles and at interfaces between dissimilar, stiff materials. Note that we aim to address mixing at the particle scale and not below. Therefore, we do not consider diffusion of material type between particles, meaning that the material of each particle remains fixed for the duration of the simulation.

We target error by exploiting three key freedoms in the SPH equations of motion presented in §1.2.2: in the choice of density estimate (§2.2.1); in the choice of free functions (§2.2.2); and in the form of the kernel function (§2.2.3). Additionally, we develop a novel method that enables the appropriate treatment of free surfaces when using these improved kernels (§2.2.4), and we use improved artificial viscosity (§2.2.5) and artificial diffusion (§2.2.6) formulations. These include new approaches both for the treatment of shocks and to weakly smooth and mitigate accumulated noise on the particle scale. We also include a term in the density evolution that re-

---

[2] SWIFT is available at `www.swiftsim.com` alongside extensive documentation and a large suite of examples.

[3] The full set of the final equations used in the REMIX scheme are listed in Appendix B. The equations of the traditional SPH scheme that we use for comparison simulations are listed in Appendix C.

ties densities to the local particle distribution (§2.2.7). These components combine into the REMIX equations of motion, given by

$$\frac{d\rho_i}{dt} = \sum_j m_j \, \frac{\rho_i}{\rho_j} \, v_{ij}^\alpha \, \frac{1}{2} \left(\left. \frac{d\widetilde{\mathcal{W}}}{dr^\alpha} \right|_{ij} - \left. \frac{d\widetilde{\mathcal{W}}}{dr^\alpha} \right|_{ji} \right) + \left(\frac{d\rho_i}{dt} \right)_{\text{difn}} + \left(\frac{d\rho_i}{dt} \right)_{\text{norm}} \,, \tag{2.1}$$

$$\frac{dv_i^\alpha}{dt} = -\sum_j m_j \, \frac{P_i + Q_{ij} + P_j + Q_{ji}}{\rho_i \, \rho_j} \, \frac{1}{2} \left(\left. \frac{d\widetilde{\mathcal{W}}}{dr^\alpha} \right|_{ij} - \left. \frac{d\widetilde{\mathcal{W}}}{dr^\alpha} \right|_{ji} \right) \,, \tag{2.2}$$

$$\frac{du_i}{dt} = \sum_j m_j \, \frac{P_i + Q_{ij}}{\rho_i \, \rho_j} \, v_{ij}^\alpha \, \frac{1}{2} \left(\left. \frac{d\widetilde{\mathcal{W}}}{dr^\alpha} \right|_{ij} - \left. \frac{d\widetilde{\mathcal{W}}}{dr^\alpha} \right|_{ji} \right) + \left(\frac{du_i}{dt} \right)_{\text{difn}} \,, \tag{2.3}$$

where $\left(d\widetilde{\mathcal{W}}/d\mathbf{r}|_{ij} - d\widetilde{\mathcal{W}}/d\mathbf{r}|_{ji} \right)/2$ are improved kernel gradient terms that are antisymmetric in the exchange of i and j for explicit conservation of momentum and energy; Q_{ij} and Q_{ji} are pairwise, artificial viscous pressures; $(d\rho_i/dt)_{\text{difn}}$ and $(du_i/dt)_{\text{difn}}$ are artificial diffusion of density and internal energy; and $(d\rho_i/dt)_{\text{norm}}$ is the kernel normalising term. Each of these are discussed in detail in their corresponding sections below.

The equations of the REMIX scheme were developed to be implemented in just three loops over particle neighbours, and without introducing any additional iterative steps. In our test simulations, performed on the COSMA8 HPC system[4], using REMIX led to a run-speed ~1.3–1.6 times longer than equivalent simulations performed with traditional SPH (and everything else unchanged). The exact amount of slowdown is problem-dependent: this range includes simulations both with and without gravity, and those using different kernel functions[5]. On the COSMA7 HPC system (which has fewer cores per node), simulations with the overhead of gravity take ~1.6–1.8 times longer, and simulations without gravity take ~2–3 times longer, depending on the test case. We find that REMIX, in addition to dealing with density discontinuities that are problematic in traditional SPH at all resolutions, is

---

[4] Simulations carried out on COSMA8 used 1 node with 128 cores and those on COSMA7 used 1 node with 28 cores. These are both part of the DiRAC cluster hosted by Durham University (`https://dirac.ac.uk/memory-intensive-durham/`).

[5] Simulations used to investigate the runtime were: 3D Kelvin–Helmholtz instabilities (§2.3.3) and planets in hydrostatic equilibrium (§2.3.9). These were tested with cubic spline and Wendland C^2 kernels.

able to achieve an improved treatment of non-discontinuous regions in simulations with over an order of magnitude lower resolution compared with equivalent traditional SPH results (§2.3.3). The effective slowdown from using REMIX is therefore much smaller in practice than the ranges above suggest, since simulations with a lower resolution (fewer SPH particles) could be used to obtain equivalent results. As such, in many cases a science simulation with REMIX would run faster than a traditional SPH simulation that would require a higher resolution to achieve a comparable level of numerical convergence. For example, the 2.9×10^5 particle REMIX Kelvin–Helmholtz instability in §2.3.3 runs over 20 times faster (on COSMA8) than the 4.7×10^6 particle traditional SPH simulation, and is closer to the converged solution[6].

2.2.1 Density estimate

In the REMIX SPH scheme we use a differential form of the density estimate: we evolve the density in time with Eqn. 2.1 rather than recalculating it each timestep (e.g. Eqn. 1.1), similarly to internal energy in traditional SPH schemes. There are three key benefits of this treatment: (1) we directly address systematic smoothing error in particle densities, which is particularly significant at density discontinuities, including those at free surfaces; (2) it allows us to constrain zeroth-order error in the equations of motion while starting from a basis of thermodynamic consistency (§2.2.2); (3) we do not require an additional loop over particle neighbours to calculate a new density each timestep. We note that particle mass is fixed throughout the simulation, so the evolution of densities is equivalent to an evolution of volumes. In §2.3.4, we show the differences in Kelvin–Helmholtz instability simulations when using the full REMIX scheme, and the REMIX scheme modified to use a traditional integral density estimate. Using our evolved density estimate, both to calculate thermodynamic quantities and in volume elements, leads to a considerable improvement in addressing spurious surface tension-like effects that suppress instability growth and mixing on the particle scale.

---

[6] See REMIX, $N = 128$ and tSPH, $N = 512$ in Fig. 2.4.

In practice, we set a density floor $\rho_{\mathrm{min,\,i}} \equiv \langle\rho_i\rangle_{\mathrm{min}} = m_i W(\mathbf{0}, h_i)$ such that $\rho_i = \rho_{\mathrm{min,\,i}}$ if the density would evolve below the minimum value. This prevents EoS extrapolation issues that arise for tiny densities in simulations involving a vacuum region.

Evolved density estimates are used frequently in SPH schemes developed for engineering applications (Antuono et al., 2010) as well as in some astrophysical SPH schemes, in particular those that include material strength models (Benz & Asphaug, 1995). However, in most astrophysical SPH schemes, an integral density estimate is preferred for its robustness: the accumulation of error in an evolved density estimate is less predictable than the relatively controlled errors in a density estimate calculated each timestep from the instantaneous local particle distribution. For instance, if left to evolve freely over many timesteps, densities could in principle take values such that volume elements m_j/ρ_j are far from normalising the kernel W_{ij}, despite the kernel being a normalised function[7]. We address these concerns with four approaches: (1) by introducing a novel term in the density evolution that re-ties densities to the local particle distribution (§2.2.7); (2) by using kernels that are normalised to the evolved densities (§2.2.3); (3) by including a weak density diffusion to smooth out accumulated noise (§2.2.6); (4) and by taking preventative measures in reducing error that could accumulate with time, reflected in the choices of our equations of motion (§2.2.2), the use of kernel functions constructed to reduce discretisation error (§2.2.3), and our improved viscosity formulation (§2.2.5).

Evolved densities are used wherever density appears in the equations of the REMIX scheme. This includes for calculating thermodynamic quantities, using the equation of state, and in all volume elements.

---

[7] Volume elements that use the interpolated density, $V_j = m_j/\langle\rho_j\rangle$, are inherently tied to kernel normalisation. The equations for kernel normalisation, Eqn. 1.12, and the integral density estimate, Eqn. 1.1, are equivalent to each other in the limit of constant density on the kernel length scale, $\langle\rho_j\rangle \to \langle\rho_i\rangle$ for all j.

2.2.2 Free functions in the equations of motion

In traditional SPH formulations, the free functions, ζ and ξ, in the equations of motions (Eqns. 1.6–1.8) typically take equal values for all particles and cancel. An alternate formulation with $\zeta = \xi = \rho$, such that the equations of motion include ratios of the densities of particles i and j, helps to constrain error in the equations of motion at density discontinuities and for irregular particle distributions on the kernel scale (Read et al., 2010). This choice avoids the use of gradients of density in the derivation of the momentum equation, by using the identity

$$\frac{\nabla P}{\rho} = \frac{\nabla P}{\rho} + \frac{P}{\rho}\nabla 1 \,, \tag{2.4}$$

rather than

$$\frac{\nabla P}{\rho} = \nabla\left(\frac{P}{\rho}\right) + \frac{P}{\rho^2}\nabla\rho \,. \tag{2.5}$$

SPH formulations using the density as the free functions have been shown to improve the treatment of mixing (Wadsley et al., 2017). For simulations using only a single ideal gas, the choice of $1/u$ as a free function is equivalent to this, with the additional assumption of constant pressure on the kernel scale (Ritchie & Thomas, 2001).

Using density as the free function in the integral form of the density estimate (Eqn. 1.1) for simulations with arbitrary EoS is not possible without iteration, since the density would be needed in the density calculation. However, using the differential form to evolve the density (Eqn. 2.1) enables us to develop the REMIX SPH scheme from a basis of full thermodynamic consistency with $\zeta_i = \rho_i$. We also use $\xi_i = \rho_i$ to reduce zeroth-order error in the momentum equation (Read et al., 2010; Wadsley et al., 2017). All densities used are the evolved densities of particles.

In §2.3.4, we demonstrate the improvements in REMIX simulations of the Kelvin–Helmholtz instability from using $\zeta_i = \xi_i = \rho_i$, compared with the REMIX scheme modified to use traditional equal-valued free functions.

2.2.3 Linear-order reproducing kernels

To reduce discretisation error, we construct kernels that explicitly satisfy the conditions given by Eqns. 1.12 and 1.13. Therefore, these kernels reproduce exact values for fields that are spatially constant or that vary linearly with position. This methodology is largely based on that of Frontiere et al. (2017). To account for spatial variations of the smoothing length, we include grad-h terms that were previously neglected. These grad-h terms take a non-standard form, compared with Hopkins (2013), since our evolved density is not tied directly to smoothing lengths through the instantaneous distribution of particles. We also modify our kernels to include a free-surface treatment (§2.2.4) to allow them to appropriately handle vacuum boundaries.

The modified kernel, $\mathcal{W}_{ij}$, is constructed so that the sum over neighbours always satisfies

$$\sum_j \mathcal{W}_{ij} V_j = 1 \ , \tag{2.6}$$

$$\sum_j \mathbf{r}_{ij} \mathcal{W}_{ij} V_j = \mathbf{0} \ . \tag{2.7}$$

We use volume elements $V_j = m_j/\rho_j$, where ρ_j are the evolved densities. We stress that for use in the equations of motion we must undergo a necessary step to make the kernel gradient terms antisymmetric in exchanges of particle pairs, to enforce the conservation of energy and momentum, as is also done by Frontiere et al. (2017). Therefore, the gradient estimates used in the equations of motion end up being not exactly first-order reproducing. Despite this, these gradient estimates show significant improvements when compared with unmodified kernels (as seen directly in §2.3.4).

To construct $\mathcal{W}_{ij}$, an unmodified SPH kernel is multiplied by a linear polynomial

$$\mathcal{W}_{ij} \equiv A_i \left(1 + B_i^\alpha r_{ij}^\alpha\right) \overline{W}_{ij} \ , \tag{2.8}$$

where $\overline{W}_{ij} \equiv [W(\mathbf{r}_{ij}, h_i) + W(\mathbf{r}_{ji}, h_j)]/2$ is a symmetrised kernel[8], and A_i and $\mathbf{B}_i$ are coefficients that satisfy Eqns. 2.6 and 2.7, as shown in Appendix A of Frontiere et al. (2017):

$$A_i = \left(\overline{m}_{0,\,i} - \left(\overline{m}_{2,\,i}^{-1}\right)^{\alpha\beta} \overline{m}_{1,\,i}^{\alpha}\, \overline{m}_{1,\,i}^{\beta} \right)^{-1} \;, \tag{2.9}$$

$$B_i^{\alpha} = -\left(\overline{m}_{2,\,i}^{-1}\right)^{\alpha\beta} \overline{m}_{1,\,i}^{\beta} \;, \tag{2.10}$$

where the geometric moments are defined as

$$\overline{m}_{0,\,i} = \sum_j \overline{W}_{ij} V_j \;, \tag{2.11}$$

$$\overline{m}_{1,\,i}^{\alpha} = \sum_j r_{ij}^{\alpha} \overline{W}_{ij} V_j \;, \tag{2.12}$$

$$\overline{m}_{2,\,i}^{\alpha\beta} = \sum_j r_{ij}^{\alpha} r_{ij}^{\beta} \overline{W}_{ij} V_j \;. \tag{2.13}$$

Greek letter indices correspond to spatial dimensions and like indices are summed over. Bars indicate the use of the symmetrised kernel in the kernel interpolation. This distinction becomes important since we use $m_{0,\,i}$, calculated similarly but using an unsymmetrised kernel, for alternative gradient estimates used later in this section and in §2.2.7.

To calculate gradient terms for the equations of motion, we require the spatial derivative of $\mathcal{W}$. We include terms that depend on the gradient of smoothing lengths, unlike Frontiere et al. (2017). We find the effects of these to be small in practice, but include them for completeness of the method – without assuming these to be negligible.

The smoothing length dependence of Eqns. 2.8–2.13 is contained within $\overline{W}_{ij}$. We therefore express the derivatives with the parameterisation

---

[8] We find this to be beneficial when we enforce the antisymmetrisation required for use in the equations of motion, as demonstrated in Appendix E. Note that for certain computational steps, this choice extends the definition of particle i's "neighbours", j, to be those that satisfy either $|\mathbf{r}_{ij}| < H_i$ or $|\mathbf{r}_{ij}| < H_j$ rather than just the first condition.

$$\mathcal{W}(\mathbf{r},\,\mathbf{r}_j) \equiv \mathcal{W}\left(\mathbf{r}-\mathbf{r}_j,\; \overline{W}(\mathbf{r}-\mathbf{r}_j, h(\mathbf{r}), h(\mathbf{r}_j)),\; A(\mathbf{r}),\; \mathbf{B}(\mathbf{r})\right) \;, \tag{2.14}$$

giving

$$\frac{d\mathcal{W}}{dr^\gamma} = AB^\gamma \overline{W} + \frac{\partial \mathcal{W}}{\partial \overline{W}}\frac{d\overline{W}}{dr^\gamma} + \frac{\partial \mathcal{W}}{\partial A}\frac{dA}{dr^\gamma} + \frac{\partial \mathcal{W}}{\partial B^\alpha}\frac{dB^\alpha}{dr^\gamma} \;. \tag{2.15}$$

When evaluated for a particle pair i, j this becomes[9]

$$\left.\frac{d\mathcal{W}}{dr^\gamma}\right|_{ij} = A_i B_i^\alpha \overline{W}_{ij} + A_i\left(1 + B_i^\alpha r_{ij}^\alpha\right)\left.\frac{d\overline{W}}{dr^\gamma}\right|_{ij} + \left(1 + B_i^\alpha r_{ij}^\alpha\right)\overline{W}_{ij}\left.\frac{dA}{dr^\gamma}\right|_i + A_i r_{ij}^\alpha \overline{W}_{ij}\left.\frac{dB^\alpha}{dr^\gamma}\right|_i \;. \tag{2.16}$$

Equations to calculate the gradients of A, B, and the geometric moments are included in Appendix B. The derivative of the symmetrised kernel is given by[10]

$$\left.\frac{d\overline{W}}{dr^\gamma}\right|_{ij} = \frac{1}{2}\left(\left.\frac{\partial W}{\partial r^\gamma}\right|_{ij} + \left.\frac{\partial W}{\partial h}\right|_{ij}\left.\frac{dh}{dr^\gamma}\right|_i - \left.\frac{\partial W}{\partial r^\gamma}\right|_{ji}\right) \;, \tag{2.17}$$

and so the inclusion of grad-h terms in the gradient calculations in practice only takes the form of the additional term in Eqn. 2.17. Both $\partial W/\partial r^\gamma$ and $\partial W/\partial h$ can be calculated directly from the kernel function (Price et al., 2018).

Finally, we require dh/dr^γ. In SPH schemes that use the traditional density estimate, dh/dr^γ do not need to be calculated explicitly (Hopkins, 2013), since smoothing lengths and densities are inherently linked. However, for the scheme presented here, where we use an evolved density estimate, we must calculate this explicitly. One approach is to directly differentiate Eqn. 1.4. However, we find that zeroth-order error from calculating grad-h terms in this way leads to spurious behaviour in

---

[9] We use the notation $\left.\frac{d\mathcal{W}}{dr^\gamma}\right|_{ij} \equiv \frac{d\mathcal{W}}{dr^\gamma}\left(\mathbf{r}_i-\mathbf{r}_j,\; \overline{W}(\mathbf{r}_i-\mathbf{r}_j, h(\mathbf{r}_i), h(\mathbf{r}_j)),\; A(\mathbf{r}_i),\; \mathbf{B}(\mathbf{r}_i)\right) \equiv \frac{d\mathcal{W}}{dr^\gamma}\left(\mathbf{r}_{ij},\; \overline{W}_{ij},\; A_i,\; \mathbf{B}_i\right)$.

[10] Note that we are taking the derivative of the continuous function $\overline{W}(\mathbf{r},\mathbf{r}_j) = [W(\mathbf{r}-\mathbf{r}_j, h(\mathbf{r})) + W(\mathbf{r}_j-\mathbf{r}, h(\mathbf{r}_j))]/2$ with respect to $\mathbf{r}$, with fixed neighbour positions $\mathbf{r}_j$, and evaluating it at $\mathbf{r}_i$. Therefore, there is only a grad-h term associated with the first term in the brackets.

simulations. We therefore calculate these by kernel interpolation. Since, in practice, $d\overline{W}/dr^\gamma$ has not been constructed yet due to the order of these operations in the loops over particle neighbours, we are unable to use these improved gradient terms for dh/dr^γ if we want to avoid introducing a 4^{th} loop. This also applies for gradient estimates in our viscosity (§2.2.5) and diffusion (§2.2.6) schemes, discussed later. We therefore require an alternative gradient estimate for these calculations. However, we must be mindful of kernel normalisation in these alternative gradient estimates, since we use evolved densities for volume elements throughout. We therefore use the kernel gradient term

$$\partial_i^\gamma \hat{W}_{ij} \equiv \frac{\partial_i^\gamma W_{ij}}{m_{0,\,i}} - \frac{W_{ij}}{m_{0,\,i}^2} \partial_i^\gamma m_0 \; , \tag{2.18}$$

where we note that the lack of bars throughout indicates the use of a standard (e.g. Wendland C^2) kernel, rather than one symmetrised by averaging with neighbouring kernels, and '∂', rather than total derivatives, indicates a lack of grad-h terms. These choices allow us to calculate these kernel gradients in two loops over particle neighbours, so they can be used here and in the artificial viscosity and diffusion schemes. Circumflexes, here and throughout, indicate the use of the normalised kernel $\hat{W}_{ij} \equiv W_{ij} \,/\, m_{0,\,i}$.

We then calculate

$$\partial_i^\gamma \hat{h} = \sum_j (h_j - h_i)\, \partial_i^\gamma \hat{W}_{ij} \frac{m_j}{\rho_j} \; , \tag{2.19}$$

and use this in place of dh/dr^γ.

All these equations combine in Eqn. 2.16 to give the gradients of the linear-order reproducing kernels. The use of these kernels reduces discretisation error in the equations of motion. In §2.3.4, we show the effect of these kernels on simulations of the Kelvin–Helmholtz instability by using either the full REMIX scheme or the REMIX scheme with unmodified, Wendland C^2 kernels.

2.2.4 Vacuum boundary treatment

We develop a method to switch the kernel gradients constructed in §2.2.3 to the unmodified spherically symmetric kernel gradients in regions identified as vacuum boundaries. We stress that this method is not applying a targeted correction to vacuum boundaries as done by, for example, Reinhardt & Stadel (2017). In fact, our evolved density estimate corrects density smoothing at discontinuous free surfaces without any need for a targeted approach. Instead, the vacuum treatment we present here is just an expansion of the form of the linear-order reproducing kernels (§2.2.3) to allow them to capture free surfaces as vacuum boundaries, a case not considered – rather than handled poorly – in their general construction.

A region with no SPH particles is not trivially equivalent to the representation of a vacuum. Since SPH particles are moving interpolation points, a region not sampled by SPH particles can be seen as analogous to a region in a grid-based code where the grid points have been removed. There is therefore no inherent information associated with these regions that would make them equivalent to a region with zero pressure, rather than a region to extrapolate into. However, if a spherically symmetric kernel, normalised to the continuum, is used to calculate pressure gradients in the equation of motion, vacuum-like behaviour is achieved. At a free surface, a particle with a spherically symmetric kernel will calculate pressure gradients equivalent to those calculated if the vacuum region were built up of particles with appropriate volumes but zero pressure[11].

This is not the case for the linear-order reproducing kernels described in §2.2.3. Since kernels are constructed to satisfy Eqns. 2.6 and 2.7 for volumes built up by particles only, the vacuum region is treated as a region to extrapolate into. SPH applications typically require the treatment of a region without SPH particles as a vacuum, or a region with negligible pressure. We therefore switch our kernel gradient terms to gradients of unmodified kernels at free surfaces:

---

[11] These gradients may not be fully equivalent in the equations of motion where the additional condition of antisymmetry in exchange of neighbours is imposed, however, they remain closely related.

$$\left.\frac{d\widetilde{\mathcal{W}}}{dr^{\gamma}}\right|_{ij} = s_i \left.\frac{d\mathcal{W}}{dr^{\gamma}}\right|_{ij} + (1 - s_i) \left.\frac{dW}{dr^{\gamma}}\right|_{ij} , \tag{2.20}$$

where s is a function that switches from 1 in regions where no vacuum boundary is detected, to 0 in regions near a vacuum boundary. Note that we smoothly switch between kernel gradients rather than the kernels themselves. This is to avoid terms with gradients of s. A switch that is accurate in identifying vacuum boundary particles *only* will inevitably have sharp spatial gradients, which could significantly influence the evolution of particles. Since we do not calculate densities by Eqn. 1.1, we do not require the direct calculation of the function whose derivative is given by Eqn. 2.20 to maintain thermodynamic consistency.

We modify the kernel gradient terms based only on parameters of the kernel function itself. Therefore, conceptually, we adapt the kernel *function* rather than making the kernel respond to the physical system simulated. For s, we use a Gaussian switch,

$$s\left(h_i|\mathbf{B}_i|\right) = \begin{cases} \exp\left[-\dfrac{(0.8 - h_i|\mathbf{B}_i|)^2}{0.08}\right] & \text{for } h_i|\mathbf{B}_i| \geq 0.8 , \\ 1 & \text{otherwise} , \end{cases} \tag{2.21}$$

where the offset, 0.8, and denominator, 0.08, of the switch are chosen empirically to identify boundary particles as those with a large $|\mathbf{B}_i|$ (Eqn. 2.10) greater than $\sim 1/h_i$. These are particles whose kernels would have to drastically change shape to deal with large anisotropies in the volume elements of particle neighbours. We find that using these values allows the switch to identify particles near free surfaces reliably without misidentifying particles in non-vacuum regions, as we show in §2.3.9, where we also demonstrate the need for this vacuum boundary treatment. In the example presented, the free surface of a Jupiter-like planet in hydrostatic equilibrium is unstable when the vacuum boundary treatment is not included. As well as its use in switching the kernel function, s is also used in the kernel normalisation term in the density evolution, as detailed in §2.2.7.

2.2.5 Artificial viscosity

Artificial viscosity is required to capture shocks in SPH simulations, whose constituent equations otherwise model adiabatic and dissipationless evolution (Monaghan & Gingold, 1983). A difficulty faced by artificial viscosity constructions is over-dissipation in regions not affected by a shock. Artificial viscosity switches, like the Balsara switch (Balsara, 1995),

$$\mathcal{B}_i = \frac{|\nabla \cdot \mathbf{v}_i|}{|\nabla \cdot \mathbf{v}_i| + |\nabla \times \mathbf{v}_i| + 0.0001\, c_i/h_i}\ , \tag{2.22}$$

where c is the sound speed, or higher-order switches like that of Read & Hayfield (2012) are used to switch artificial viscosity off in shearing regions. Time-dependent viscosity parameters have also been developed (Borrow et al., 2022; Cullen & Dehnen, 2010; Morris & Monaghan, 1997) to reduce over-dissipation.

Recently, the limiting of artificial viscosity by the use of reconstructed velocities at particle-pair midpoints has been demonstrated to be an effective alternative approach (Frontiere et al., 2017; Pearl et al., 2022; Rosswog, 2020). For each particle pair, two velocities are estimated at the midpoint of the pair based on Taylor expansions from each particle using their individual velocities and estimated velocity gradients. The difference between these velocities is then used in the viscosity scheme instead of the relative velocity of the particles themselves. This is the approach taken in REMIX. We use linear reconstruction as we find further improvements due to quadratic reconstruction to be small, as also noted by Rosswog (2020). If the velocity field is locally linear, artificial viscosity would effectively be switched off with linear reconstruction. For schemes that use linear reconstruction, the viscosity in shearing regions where the velocity field is not exactly locally linear is not negligible and will still influence the fluid behaviour. However, this results in a helpful effect, acting as a weak artificial diffusion of momentum that smooths particle noise in the velocity field by guiding it towards being locally linear on the particle scale.

Our artificial viscosity treatment is largely based on those of Frontiere et al. (2017) and Rosswog (2020), with some additional, novel approaches. As detailed below, a slope limiter is used to prevent reconstruction at discontinuities, thereby increasing artificial viscosity where it is required for shock capturing. However, we find that a slope limiter alone does not effectively switch off reconstruction, because the velocity gradients used to construct it are inherently smoothed by their calculation using a smoothing kernel. Therefore, they do not identify sharp discontinuities well. We introduce a Balsara switch (Eqn. 2.22) into the slope limiter term to switch off reconstruction at shocks more effectively. Here we calculate $|\nabla \cdot \mathbf{v}_i|$ and $|\nabla \times \mathbf{v}_i|$ in the Balsara switch using the kernel gradient term given by Eqn. 2.18, and also use these same gradient estimates for the velocity gradients used in the linear reconstruction,

$$\partial_i^\gamma \hat{v}^\alpha = \sum_j (v_j^\alpha - v_i^\alpha)\, \partial_i^\gamma \hat{W}_{ij} \frac{m_j}{\rho_j} \ . \tag{2.23}$$

The velocity reconstructed to the midpoint of a particle pair is given by

$$\tilde{v}_{ij}^\alpha = v_i^\alpha + \frac{1}{2}\left(1 - \mathcal{B}_i^{\mathrm{SL}}\right) \Phi_{v,\,ij} \left(r_j^\gamma - r_i^\gamma\right) \partial_i^\gamma \hat{v}^\alpha \ , \tag{2.24}$$

where $\mathcal{B}_i^{\mathrm{SL}}$ is the standard Balsara switch (Eqn. 2.22), and the SL (slope limiter) superscript just indicates its use in conjunction with the slope limiter. Φ_{ij} is the van Leer slope limiter (Van Leer, 1974), given by

$$\Phi_{ij} = \begin{cases} 0 & \text{for } A_{ij} < 0 \ , \\ \dfrac{4A_{ij}}{(1+A_{ij})^2} \exp\left[-\left(\dfrac{\eta_{ij}^{\min} - \eta_{\mathrm{crit}}}{0.2}\right)^2\right] & \text{for } \eta_{ij}^{\min} < \eta_{\mathrm{crit}} \ , \\ \dfrac{4A_{ij}}{(1+A_{ij})^2} & \text{otherwise} \ , \end{cases} \tag{2.25}$$

where the additional Gaussian term in Eqn. 2.25 switches the slope limiter to 0 for particle pairs with a small separation. $\eta_{ij}^{\min}$ is the smaller value of $|\boldsymbol{\eta}_{ij}|$ and $|\boldsymbol{\eta}_{ji}|$, where $\boldsymbol{\eta}_{ij} = (\mathbf{r}_i - \mathbf{r}_j)/h_i$ and similarly for the exchanged particle indices. η_{crit} represents a separation closer than one would expect from the distribution of the rest of the particle's neighbours. For viscosity calculations, we use the ratio of projected velocity gradients $A_{ij} \equiv A_{v,\,ij}$ given by

$$A_{v,\,ij} = \frac{\partial_i^\alpha \hat{v}^\beta (\mathbf{r}_j - \mathbf{r}_i)^\beta (\mathbf{r}_j - \mathbf{r}_i)^\alpha}{\partial_j^\gamma \hat{v}^\phi (\mathbf{r}_j - \mathbf{r}_i)^\gamma (\mathbf{r}_j - \mathbf{r}_i)^\phi} \, . \tag{2.26}$$

For η_crit we use

$$\eta_\text{crit} = \frac{1}{h_i} \left(\frac{1}{\sum_j W_{ij}} \right)^{1/d} \equiv \frac{1}{\eta_\text{kernel}} \, , \tag{2.27}$$

where the equivalency is due to the definition of the smoothing length in Eqn. 1.4. Note that the term in brackets is an approximation of the particle volume assuming neighbours with equal volumes.

The reconstructed velocities appear in the artificial viscosity formulation through

$$\mu_{ij} = \begin{cases} \dfrac{\tilde{\mathbf{v}}_{ij} \cdot \boldsymbol{\eta}_{ij}}{\boldsymbol{\eta}_{ij} \cdot \boldsymbol{\eta}_{ij} + \epsilon^2} & \text{for } \tilde{\mathbf{v}}_{ij} \cdot \boldsymbol{\eta}_{ij} < 0 \, , \\ 0 & \text{otherwise} \, , \end{cases} \tag{2.28}$$

and similarly for μ_{ji} with all particle indices exchanged throughout the calculations. $\epsilon = 0.1$ is a small constant. Similarly to the artificial viscosity of Monaghan & Gingold (1983), each pressure term in the equations of motion is modified with the addition of a pairwise viscous pressure[12]. The viscous pressure terms Q_{ij} combine a linear bulk viscosity term and a quadratic Von Neumann–Richtmyer viscosity term (VonNeumann & Richtmyer, 1950),

$$Q_{ij} = \frac{1}{2} \left(a_\text{visc} + b_\text{visc} \mathcal{B}_i^\text{visc} \right) \rho_i \left(-\alpha c_i \mu_{ij} + \beta \mu_{ij}^2 \right) \, , \tag{2.29}$$

and similarly for Q_{ji} with all particle indices exchanged throughout the calculations. The constants α and β set the strengths of the bulk and Von Neumann–Richtmyer terms. The constants a_visc and b_visc set the strength of the viscosity in regions of different flow, based on the Balsara switch, $\mathcal{B}_i^\text{visc}$.

The REMIX artificial viscosity scheme differs from those of Frontiere et al. (2017) and Rosswog (2020) in some notable aspects: firstly, the Balsara switch, $\mathcal{B}_i^\text{SL}$, is included in the slope limiter term (in Eqn. 2.24). This avoids reducing the artificial

---

[12] P_i becomes $P_i + Q_{ij}$ and P_j becomes $P_j + Q_{ji}$.

viscosity where it is needed, leading to a more effective targeting of shocks. This allows us to introduce a factor of $1/2$ in Eqn. 2.29 to recover equations more closely equivalent to those in Price (2012). Otherwise, the contributions from both Q_{ij} and Q_{ji} would effectively lead to this being a factor of 2 stronger, which is to some extent mitigated by those schemes being ineffective at switching off velocity reconstruction in shocks. Secondly, we use $\alpha = 1.5$ and $\beta = 3$ as we find that these slightly larger constants, compared with $\alpha = 1$ and $\beta = 2$ as used by Frontiere et al. (2017) and Rosswog (2020), help to dissipate spurious oscillations in shocks in 3D. This is consistent with typical values used in planetary impact simulations (e.g. Canup, 2004; Reinhardt & Stadel, 2017). Thirdly, we use an additional Balsara switch directly in Eqn. 2.29, which, combined with the values we use for $a_{\rm visc} = 2/3$ and $b_{\rm visc} = 1/3$, acts to switch between $\alpha = 1.5$ and $\beta = 3$ in shocks and $\alpha = 1$ and $\beta = 2$ in shearing regions. Here we make relatively conservative choices to limit the effect of artificial viscosity in smoothing particle noise in shearing regions, despite finding it to be a useful effect, owing to the velocity reconstruction to particle midpoints. Our artificial viscosity scheme is constructed to be less dissipative in shearing regions and to target shocks more effectively than similar schemes. These choices are all discussed in more detail in Appendix F.

2.2.6 Artificial diffusion

Artificial diffusion of internal energy, or "artificial conduction[13]", is frequently used to smooth accumulated noise in particle internal energies (Monaghan, 1997) or entropies (Schaller et al., 2015), and to improve the treatment of density discontinuities in ideal gas-only simulations (Price, 2008). As with artificial viscosity, a targeted approach is desirable to avoid artificial conduction playing a dominant role in the thermodynamic evolution, instead of acting as a correction on the particle scale (Read & Hayfield, 2012).

In some SPH schemes, relatively strong artificial conduction is used to address

---

[13] In later sections, we use "artificial diffusion" to refer to cases that include the diffusion of both density and internal energy and "artificial conduction" where there is only diffusion of internal energy.

kernel smoothing at density discontinuities by smoothing particle internal energies over kernel length scales (Price et al., 2018). For a single equation of state, with no phase transitions, this leads to a smooth pressure field in the continuous limit. However, this is not an appropriate treatment in simulations with multiple and/or complex materials, where smooth density and internal energy fields do not necessarily lead to smooth pressures. Additionally, even in ideal gas-only simulations, this does not completely solve the issue, since (1) artificial conduction becomes a less effective correction at large density discontinuities; (2) in simulations with gravity, strong diffusion will disturb a system's hydrostatic equilibrium; (3) artificial conduction does not attempt to address the source of kernel smoothing error directly, instead it alters the physical system itself to one without discontinuities.

In simulations that use an evolved density estimate (Eqn. 1.6), a similar artificial diffusion term can be used in the evolution of densities, for example, in the δ-SPH formulation, used predominantly for engineering applications (Antuono et al., 2010, 2012; Sun et al., 2021).

In REMIX, we include artificial diffusion of specific internal energy and of density, both to improve the treatment of shocks and to smooth accumulated noise on the particle scale, using reconstruction to particle midpoints (Antuono et al., 2010; Rosswog, 2020). Similarly to the phase dependence in the diffusion schemes of Sun et al. (2021) and Pearl et al. (2022), we only allow diffusion between particles of the same material type. Without this distinction, artificial diffusion of internal energy between different materials would cause unphysical evolution, since smoothing would be based on *internal energy* and not *temperature*. Diffusing density between different materials would lead to density discontinuities at material interfaces returning to a similar, smoothed state as in simulations with smoothing error in the density estimate.

The diffusion terms in the equations of motion take the form

$$\left(\frac{du_i}{dt}\right)_{\rm difn} = \sum_j \kappa_{ij}\left(a_u + b_u\mathcal{B}_{ij}^{\rm difn}\right) v_{{\rm sig},ij}\left(\tilde{u}_j - \tilde{u}_i\right)\frac{m_j}{\rho_{ij}}\frac{1}{2}\left|\left.\frac{d\widetilde{\mathcal{W}}}{d\mathbf{r}}\right|_{ij} - \left.\frac{d\widetilde{\mathcal{W}}}{d\mathbf{r}}\right|_{ji}\right| , \tag{2.30}$$

$$\left(\frac{d\rho_i}{dt}\right)_{\rm difn} = \sum_j \kappa_{ij}\left(a_\rho + b_\rho\mathcal{B}_{ij}^{\rm difn}\right) v_{{\rm sig},ij}\left(\tilde{\rho}_j - \tilde{\rho}_i\right)\frac{\rho_i}{\rho_j}\frac{m_j}{\rho_{ij}}\frac{1}{2}\left|\left.\frac{d\widetilde{\mathcal{W}}}{d\mathbf{r}}\right|_{ij} - \left.\frac{d\widetilde{\mathcal{W}}}{d\mathbf{r}}\right|_{ji}\right| , \tag{2.31}$$

where $\kappa_{ij} = 1$ for particles of the same material and $\kappa_{ij} = 0$ otherwise. The average Balsara switch for each particle pair is used, $\mathcal{B}_{ij}^{\rm difn} = (\mathcal{B}_i + \mathcal{B}_j)/2$, for conservation. We take the signal velocity to be $v_{{\rm sig},\,ij} = |\tilde{\mathbf{v}}_i - \tilde{\mathbf{v}}_j|$ and do not draw any distinctions between simulations with and without gravity (unlike some previous works (Price et al., 2018; Rosswog, 2020)), since we aim to validate the full REMIX formulation independently of specific simulation properties. The parameters a_u and a_ρ set the strength of the artificial conduction in shearing regions (where $\mathcal{B}_{ij}^{\rm difn} \to 0$) and are increased to $a_u + b_u$ and $a_\rho + b_\rho$ in shocks. In shearing regions we choose to have low amounts of diffusion to avoid this strongly influencing thermodynamic evolution, and to allow for persisting and emergent discontinuities. We therefore use $a_u = a_\rho = 0.05$, similarly to Rosswog (2020). In the presence of shocks we find that we need a much larger amount of diffusion to prevent spikes in density and internal energy, and so we use $b_u = b_\rho = 0.95$. We motivate and test the sensitivity of these choices in Appendix F. The volume elements in Eqn. 2.30 are chosen to conserve energy. In Eqn. 2.31, they include an additional ratio of densities, to conserve volume in each pairwise interaction[14]. Although conserving volume in a pairwise interaction between particles is not strictly necessary, we find that it improves the treatment of the density diffusion in shocks.

When calculating the artificial diffusion terms, internal energies and densities are reconstructed to particle midpoints similarly to the velocities in the artificial viscosity scheme via

$$\tilde{u}_i = u_i + \frac{1}{2}\Phi_{u,\,ij}\left(r_j^\gamma - r_i^\gamma\right)\partial_{\kappa,\,i}^\gamma\,\hat{u}\ , \tag{2.32}$$

---

[14] Substituting $\frac{d\rho_i}{dt} = -\frac{\rho_i^2}{m_i}\frac{dV_i}{dt}$ and solving for $\frac{dV_i}{dt}$ gives an equation antisymmetric in the exchange of particles.

$$\tilde{\rho}_i = \rho_i + \frac{1}{2}\Phi_{\rho,\,ij}\left(r_j^\gamma - r_i^\gamma\right)\partial_{\kappa,\,i}^\gamma\,\hat{\rho}\ . \tag{2.33}$$

The derivatives are calculated using only particles of the same material species as

$$\partial_{\kappa,\,i}^\gamma\,\hat{u} = \sum_j \kappa_{ij}\left(u_j - u_i\right)\partial_i^\gamma \hat{W}_{ij}\frac{m_j}{\rho_j}\ , \tag{2.34}$$

$$\partial_{\kappa,\,i}^\gamma\,\hat{\rho} = \sum_j \kappa_{ij}\left(\rho_j - \rho_i\right)\partial_i^\gamma \hat{W}_{ij}\frac{m_j}{\rho_j}\ . \tag{2.35}$$

The material dependence of these gradients helps to preserve real discontinuities at material boundaries.

The slope limiter is calculated in the same way as for the viscosity, Eqn. 2.25, but with $A_{ij} = A_{u,\,ij}$ and $A_{ij} = A_{\rho,\,ij}$ given by

$$A_{u,\,ij} = \frac{\partial_{\kappa,\,i}^\alpha\,\hat{u}(\mathbf{r}_j - \mathbf{r}_i)^\alpha}{\partial_{\kappa,\,j}^\beta\,\hat{u}(\mathbf{r}_j - \mathbf{r}_i)^\beta}\ , \tag{2.36}$$

$$A_{\rho,\,ij} = \frac{\partial_{\kappa,\,i}^\alpha\,\hat{\rho}(\mathbf{r}_j - \mathbf{r}_i)^\alpha}{\partial_{\kappa,\,j}^\beta\,\hat{\rho}(\mathbf{r}_j - \mathbf{r}_i)^\beta}\ . \tag{2.37}$$

Although our diffusion scheme technically includes material dependence, this is not a correction targeted at material boundaries, nor with any dependence on the actual EoS. Rather, we actively turn off these parts of our method for particles of different species. Our artificial diffusion scheme is used to improve the treatment of shocks, and to weakly smooth accumulated noise. It is not used to address surface tension-like effects that prevent mixing and instability growth at density discontinuities, even in our ideal gas-only simulations.

2.2.7 Normalising term

We add a normalising term to the density evolution equation. This aims to evolve densities to reflect the distribution of mass in nearby particles, particularly in regions where particle volume elements systematically fail to satisfy the normalisation of the kernel. Since error accumulates in the evolution of densities based on timescales set by the divergence operator used in the equations of motion, we set the normalising term to act over timescales determined by the motion of particles. This also allows particles to move in response to changes in density caused by the normalising term.

Particle volume elements should approximately satisfy $\sum_j W_{ij} V_j = 1$ (Eqn. 1.12) for a normalised kernel function. However, this condition will not be satisfied either if particle densities are poor estimates of the underlying field or if particle masses do not appropriately represent the mass distribution of the fluid. Our methods inherently conserve mass, as particle masses do not evolve during the simulation, and are fully Lagrangian. Therefore, we choose to maintain the simplicity and computational stability of this construction, and address discrepancies in volume elements through particle densities rather than through particle masses or their distribution. We do this by including an additional term in the density evolution, which we refer to as the "normalising term", that evolves densities towards a set of volume elements that aim to appropriately build up the continuous simulation volume. We note that the role of this term is not to obtain volume elements that exactly satisfy normalisation for all particle kernels at any given time, but rather to keep volume elements loosely tied to kernel normalisation and to address regions with systematic discrepancies.

To construct our normalising term, we consider the zeroth geometric moment of the unmodified kernel,

$$m_{0,i} = \sum_j W_{ij} V_j \ , \tag{2.38}$$

where $m_{0,i} = 1$ if the kernel W_{ij} is normalised over the volume elements $V_j = m_j / \rho_j$. For a single particle i, we could trivially satisfy this condition by modifying the density of the particle and all its neighbours, j, by replacing ρ_j with $m_{0,i}\, \rho_j$. However, this does not imply that $m_{0,j} = 1$ for all j, which will all have different $m_{0,j}$ and different sets of neighbours. But if there are systematic discrepancies in m_0 for many neighbouring particles, then modifying densities in a similar way for all these particles *will* move them closer to $m_{0,j} = 1$. For instance, consider a region where particles have systematically too low density, leading to a local trend of

$m_{0,\,j} > 1$. Here, increasing the densities will evolve these particles towards $m_{0,\,j} = 1$ and towards a density field that better represents the local mass distribution. In practice, we capture this behaviour with a smooth evolution in time. Unlike in the initial naïve example of modifying the densities of all j to satisfy $m_{0,\,i} = 1$ for i only, we evolve the density of i only, based on its own $m_{0,\,i}$. This reduces the risk of emergent chaotic behaviour and still captures the desired behaviour in regions of systematic trends away from kernel normalisation. The normalising term in the density evolution equation takes the form

$$\left(\frac{d\rho_i}{dt}\right)_{\text{norm}} = \alpha_{\text{norm}}\, s_i\, (m_{0,\,i} - 1)\, \rho_i \sum_j v_{\text{norm},\, ij}\, \frac{m_j}{\rho_{ij}} \frac{1}{2} \left| \left.\frac{d\widetilde{\mathcal{W}}}{d\mathbf{r}}\right|_{ij} - \left.\frac{d\widetilde{\mathcal{W}}}{d\mathbf{r}}\right|_{ji} \right| . \tag{2.39}$$

where $\alpha_{\text{norm}} = 1$ is a constant and $v_{\text{norm},\, ij} = |\mathbf{v}_i - \mathbf{v}_j|$ is the effective signal velocity. Eqn. 2.39 aims to contribute to a weak evolution of ρ_i towards $m_{0,\,i}\, \rho_i$. We include the vacuum switch, s, described in §2.2.4, since the kernel should not be normalised by particle volume elements at vacuum boundaries[15]. Here, we use the same volume elements and kernel gradient terms as are used in the diffusion of internal energy (Eqn. 2.30), despite not being motivated by conservation in this term, since it does not represent the exchange of a quantity between particles. We use these so that the timescale of the normalising evolution is based on terms in the sum that are equal for both particles in each pairwise interaction. This prevents individual particles dominating in the corrective evolution. Using a timescale that depends on particle motion $v_{\text{norm},\, ij}$ rather than, for example the sound speed, allows particles to react and move in response to changes in density caused by the normalising term. We find that using an effective signal velocity that depends on the sound speed, even with a small multiplication factor, can lead to spatial oscillations in density, because densities change to attempt to satisfy normalisation faster than particles can respond to these changes.

We show the effect of this term in simulations in §2.3.3 and §2.3.9. In particular, we show that without this term, an example Jupiter-like planet in hydrostatic equi-

---

[15] At vacuum boundaries, one would instead expect $m_{0,\,i} \approx 1/2$.

librium will develop numerical instabilities as particles with low evolved densities, but are in regions of high particle number density, move from the planet's surface towards its core (§2.3.9). In less extreme cases, the normalising term does not have a significant effect on hydrodynamics, although it does lead to particle densities that are generally closer to satisfying kernel normalisation, $m_{0,\,i} = 1$.

2.3 Hydrodynamic tests

In this section, we validate REMIX in simulations to test its ability to capture physically realistic fluid behaviour. The primary tests are performed with particles of equal mass across the simulation, though we also include a subset of additional simulations for direct comparisons with past work, where particles are placed onto a regular grid but have different masses. We refer to these two cases as "equal mass" and "equal spacing" throughout the following sections. The choice to focus on simulations with equal mass is made to validate our methods for science applications where particle densities and configurations can evolve significantly from their initial states, so particle masses cannot be easily chosen in the initial configuration to address errors. All simulations are performed in 3D, to account for effects that do not change predictably when increasing the number of dimensions, such as due to more freedom in particle configurations, or the change in scaling between neighbour number and length scale of particle interactions[16]. Additionally, in figures showing simulation snapshots, we deliberately plot individual particles rather than the smooth, reconstructed fields shown in some works. It is particularly important to visualise small-scale behaviour of simulations that aim to improve the treatment of density discontinuities where the effects that suppress mixing act on the particle scale.

We present results for the following hydrodynamic test scenarios:

- the square test (§2.3.1), where we investigate the treatment of density discon-

---

[16] A 2D simulation will have a lower number of neighbours for a given smoothing length than the equivalent 3D simulation. Increasing η_{kernel} to compensate for this would lead to kernel smoothing over a larger length scale.

tinuities in static equilibrium;

- the Sod shock tube (§2.3.2), where we investigate the treatment of shocks;
- the Kelvin–Helmholtz instability both with an ideal gas EoS (§2.3.3) and between different, stiff materials set up to be representative of iron & rock material boundaries in an Earth-like planet (§2.3.4);
- the Rayleigh–Taylor instability, also both with an ideal gas EoS (§2.3.5) and with iron & rock (§2.3.6);
- the blob test (§2.3.7), with which we investigate the onset of turbulence due to unseeded instabilities in both subsonic and supersonic regimes;
- the Evrard collapse (§2.3.8), which is used to test the interaction of our hydrodynamic treatment with gravity and shocks;
- and finally, planets in hydrostatic equilibrium (§2.3.9), which we consider as a test scenario that combines gravity, complex-material boundaries, and a vacuum boundary.

The initial conditions needed to perform these tests are included as examples in the open-source SWIFT code.

We include comparisons with simulations carried out both using a traditional SPH formulation ("tSPH") and a traditional formulation that includes artificial conduction of internal energy ("tSPH + cond."), with full details in Appendix C. These are used to demonstrate the motivation and need for many of the improvements in REMIX. We note that in most ideal gas tests, we follow the convention of past work and leave quantities unitless.

2.3.1 Square test

The "square test" is used to investigate spurious surface tension-like effects from sharp discontinuities in a system that should be in static equilibrium (Saitoh & Makino, 2013). Here we test both an equal spacing scenario, i.e., with different particle masses in the two regions, and an equal mass scenario. The significant contributions from both smoothing and discretisation error (§1.3.1, §1.3.2) at the density discontinuity makes the equal mass test particularly challenging for SPH.

A square (or cube) of fluid of higher density is initiated in pressure equilibrium with the surrounding region of low density fluid. Since the fluid experiences no gradients in pressure, other than those created by numerical errors, the shape of the square should not distort with time. In tSPH simulations, spurious surface tension-like effects at the density discontinuity leads to non-zero accelerations and a deformation of the square (Springel, 2010a). Typically, this test is carried out in 2D however, here we simulate a more challenging 3D cube with its effectively "sharper" higher-dimension corners, similarly to Rosswog (2020).

First, for the equal spacing scenario, we use initial conditions set up to match those of Rosswog (2020). 40^3 particles are placed in a simple cubic lattice with spacing 1/40 between adjacent particles. The simulation box is periodic and has length 1 in each of the x, y, z directions. Masses are chosen such that $m_i = \rho(\mathbf{r}_i)/40^3$, with densities $\rho = 4$ in the region $-0.25 < x, y, z < 0.25$ and $\rho = 1$ otherwise. An ideal gas EoS with $\gamma = 5/3$ is used for all particles. Initial internal energies are set to give a uniform pressure[17] of $P(\rho, u) = 2.5$.

The evolution of the equal spacing square test, carried out using each of tSPH; tSPH with artificial conduction; and REMIX, is shown in the top panels of Fig. 2.1. In the equal spacing scenario, the major contribution to spurious surface tension is due to the smoothing of the density field. The contribution of discretisation error is small, due to the well-ordered particle distribution and use of a relatively high-order kernel. With tSPH, the cube quickly deforms to a more stable, spherical shape, as illustrated by the upper, top-left panels of Fig. 2.1. Artificial conduction reduces the effects of smoothing error and so a square shape persists for longer, although the sharpness of the discontinuity is not maintained (Fig. 2.1 lower, top-left panels). With REMIX, particle motion is negligible, relative to the particle separation, and the cube retains its shape (Fig. 2.1 top-right panels). This is in large part due

---

[17] We note the use of the unsmoothed density ρ rather than the smoothed $\langle\rho\rangle$ used to set the internal energies of the initial conditions. Therefore tSPH simulations are not initialised in pressure equilibrium, due to smoothing error in the density estimate.

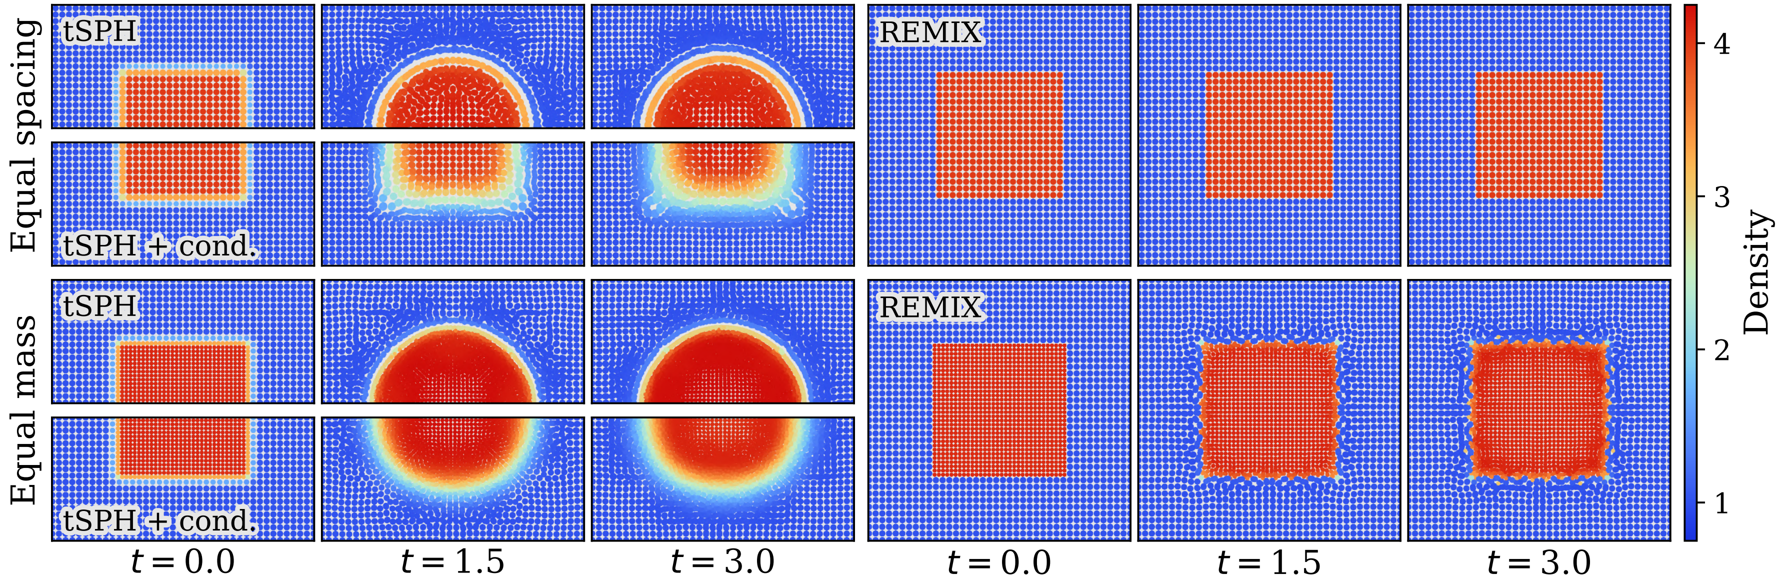

Figure 2.1: Central cross-sections from 3D “square test” simulations. Snapshots of the cube are shown at three times from simulations with equal initial particle spacing (top full-row) and equal particle mass across the simulation (bottom full-row). Simulations were carried out using three SPH formulations: tSPH, tSPH with artificial conduction, and REMIX. Individual particles are plotted at their positions in the x–y plane on a grey background, and coloured by their density. Particles in a slice of thickness 0.1 are plotted, so the grey background is visible in regions that have maintained their grid alignment in z from the initial conditions.

to the use of the evolved density estimate, which prevents density smoothing – and therefore spurious pressures – at the discontinuities. Our choice of the free functions in the equations of motion and kernel construction also help in reducing discretisation error to achieve these results.

Next, we consider the more challenging case for SPH: the use of equal mass particles, which leads to particles set up in considerably different grid-spacings interacting at the density discontinuity. Particles in the low density region are placed in the same configuration as in the equal spacing scenario. Then, instead of increasing particle masses in the high density region, the particle spacing is decreased and masses are kept the same as in the low density region. To satisfy these conditions while closely matching the density ratio in the equal spacing test, the high density region is given a grid-spacing of a factor 0.625 finer than the grid-spacing of the low density region. This corresponds to a density of 4.096. The new spacing of high density particles is chosen such that the layers of particles on either side of discontinuities are separated by the mean of the two grid-spacings, for all cube faces.

The evolution of this square test with equal mass particles is shown in the bottom panels of Fig. 2.1. There is now a large contribution of both smoothing and discretisation error in both of the traditional SPH formalisms. As such, the cube quickly deforms, even with conduction acting to reduce smoothing error. In the REMIX formulation, some minor deformation can be observed over these timescales. However, the general shape is maintained (Fig. 2.1 bottom-right panels). We note that although past work typically shows 2D square test evolution over longer timescales than those of our plotted snapshots, our plots show times later than the comparable 3D tests in Rosswog (2020), beyond the time at which their equal spacing cubes have deformed. Reducing the effects of artificial surface tension requires all of (1) a density estimate that does not smooth density discontinuities, (2) our choice of equations of motion, and (3) improved gradient estimates. In the REMIX simulation, artificial diffusion is not the dominant source of correction, as discontinuities in both density and internal energy remain sharp.

If the linear-order reproducing kernels are used in the equations of motion with-

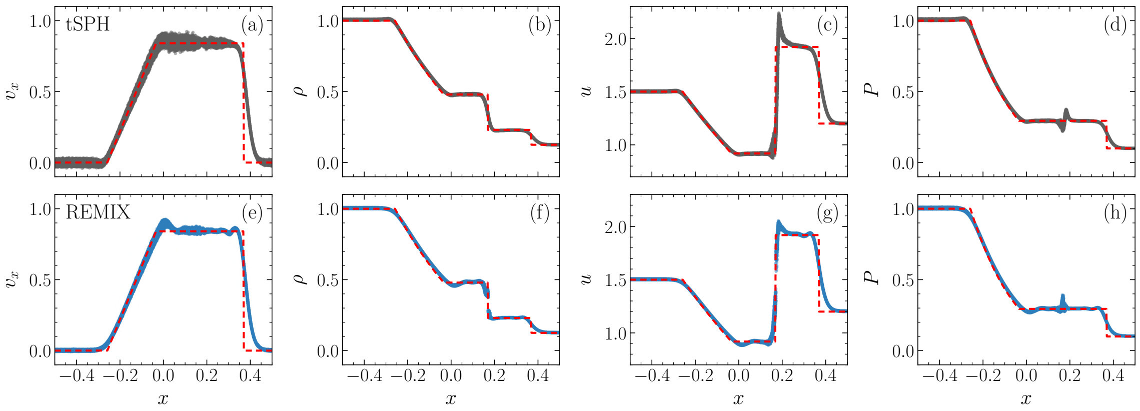

Figure 2.2: 3D Sod shock tube at time $t = 0.2$ simulated using tSPH and REMIX. Plots show velocity in the x-direction, v_x, density, ρ, specific internal energy, u, and pressure, P, of individual particles plotted against their x-position for tSPH (a–d) and REMIX (e–h) respectively. The red, dashed line shows a reference solution, solved for directly by using a Riemann solver. All particles are plotted.

out the antisymmetrisation, which is needed to enforce conservation, the square will remain undisturbed over much longer timescales, even in the equal mass case. The difference in outcome between using the conservative, antisymmetric construction and the exactly linear reproducing construction is sensitive to the kernel function used to construct the linear reproducing kernel. Therefore, reducing the additional error introduced in antisymmetrisation becomes an important consideration when choosing the form of the kernel from which the linear-order reproducing kernels are constructed. This can be seen in Appendix E, where we present sensitivities in these results to different elements of the REMIX construction.

2.3.2 Sod shock tube

The "Sod shock tube" (Sod, 1978) is used to assess the shock capturing capabilities of our hydrodynamic scheme. This is a classic Riemann problem with a known analytic solution. Since the inclusion of artificial viscosity and diffusion are necessary to deal with shocks in the REMIX scheme, we also use this test to motivate choices made in the artificial viscosity and diffusion formulations, as detailed in Appendix F. The choices made in the viscosity scheme relating to this test focus on reducing ringing oscillations behind the shock. The diffusion scheme focuses on reducing the size of spikes in density and internal energy at the discontinuity.

Ideal gas, $\gamma = 5/3$, particles of equal mass are placed in a periodic 3D domain with size 2 in each of x, y, z directions, centred at (0, 0, 0). We use two glass configurations, scaled appropriately for the two regions of different initial density: $\rho_1 = 1$ in the region $x < 0$ and $\rho_2 = 1/8$ in the region $x > 0$. Initial internal energies are set such that $P_1 = 1$ and $P_2 = 0.1$. Simulations have a total of 589,824 equal mass particles.

The particle velocities in the x-direction, densities, internal energies, and pressures at a time $t = 0.2$ are shown in Fig. 2.2. The shock is captured well with REMIX, and the particles align with the reference solution. Noise in particle velocities is reduced compared with tSPH. The size of the spike in internal energy is also reduced. The pressure blip could be further smoothed by increasing the strength of our artificial diffusion scheme, through choices of the a and b factors. However, we choose to take a conservative approach to artificial diffusion to avoid deviating far from the thermodynamically consistent core equations of motion.

2.3.3 Kelvin–Helmholtz instability – ideal gas

The Kelvin–Helmholtz instability (KHI) is the first test we use to investigate the treatment of mixing and dynamic instability growth in our simulations. The KHI arises at shearing interfaces in fluids (Chandrasekhar, 1961). Perturbations at the interface grow to form vortices that act to cascade energy to shorter length scales. As such, the KHI plays a significant role in the onset of turbulence in physical systems. Capturing the growth of the KHI has therefore been widely adopted as a benchmarking test to assess a numerical method's ability to simulate turbulence-driven mixing, as well as mixing on the particle scale. However, unlike the other tests above, an analytical solution does not exist for the KHI.

Here we first consider the growth of these instabilities at shearing density contrasts in an ideal gas. All simulations presented are carried out in 3D, with a thin z direction depth relative to the other dimensions, similarly to Hopkins (2013), Read et al. (2010), and Rosswog (2020). We focus primarily on cases with a sharp density discontinuity and equal mass particles. This is in contrast with an alternative setup with which we directly compare our results with a reference solution (McNally et al., 2012), where we consider an initially smoothed discontinuity and equal particle spacing. Although the use of this second form of initial conditions with smooth initial densities and velocities is motivated by the existence of a converged solution, these choices change the physical system to one with inherently less smoothing and discretisation error, which are the main effects of interest that normally suppress instability growth in SPH simulations. These smooth initial conditions therefore do not give the full picture of an SPH scheme's ability to capture KHI growth at sharp density discontinuities, where these sources of error can play a dominant role. This is particularly important at material boundaries, where smoothing the density discontinuity between different materials may lead to particles of both materials occupying extreme regions of their EoS phase space, so considering deliberately smoothed, equilibrium initial conditions would not be representative of a physical system.

Traditional formulations of SPH struggle to capture the KHI (Agertz et al., 2007), with the growth of the instability being strongly suppressed. In particular, for shearing density discontinuities, smoothing in the density estimate leads to surface tension-like effects that act to artificially stabilise the interface. Additionally, for simulations where SPH particles in both density regions have equal mass, or configurations that give similarly anisotropic local particle distributions at the interface, leading-order error in the momentum equation will also contribute significantly to this spurious surface tension-like effect. Not only do these effects act to suppress mixing by hampering the large-scale evolution of naturally arising instabilities that should act to drive mixing, but they will also impede particles crossing interfaces, thereby suppressing mixing both indirectly and directly.

The growth of a mode of wavelength λ is characterised by the timescale (Price, 2008; Rosswog, 2020)

$$\tau_{\mathrm{KH}} = \frac{(\rho_1 + \rho_2)\ \lambda}{\sqrt{\rho_1 \rho_2}\,|v_1 - v_2|}\ , \tag{2.40}$$

where ρ_1 and ρ_2 are the densities in regions separated by the shearing interface and $|v_1 - v_2|$ is their relative speed. We use this parameterisation so that comparisons can be drawn at the same τ_{KH} between simulations with different initial conditions, since we consider KHIs with both smoothed and sharp interfaces, for different density ratios, and between different materials. We note that initial conditions with and without initial smoothing of fields at the interface are physically different systems, so we do not expect converged results between the two.

In the absence of stabilising influences such as physical surface tension or gravity, a shearing discontinuity is unstable to perturbation modes of all wavenumbers (Chandrasekhar, 1961). In a realistic system satisfying these conditions, instability will always be triggered, as even the smallest local inhomogeneities will seed mode growth. Similarly, in a simulation, numerical error will inevitably trigger instability at shearing discontinuities. The wavenumbers of error-seeded modes are sensitive not only to the numerical methods used and the construction of initial conditions, but also to the resolution of the simulation: a higher resolution simulation will be able to resolve the excitation of a wider range of mode wavelengths (Robertson et al., 2010). The growth of KHIs at sharp discontinuities can therefore not be used reliably for convergence studies.

McNally et al. (2012) and Robertson et al. (2010) construct KHI initial conditions with smooth initial velocities and densities across the shearing interface. They show that the inclusion of a well-resolved transition region acts to stabilise the system, suppressing modes other than those deliberately seeded in the initial conditions. They demonstrate convergence and present a well-posed method to benchmark the early evolution of KHI simulations. In §2.3.3 below, we present REMIX simulations using the initial conditions of McNally et al. (2012), including quantitative comparisons of mode growth with their converged reference solution. In §2.3.3, we present KHI simulations with sharp discontinuities in density and velocity across the interface. Although we cannot make quantitative comparisons of this more challenging case with converged reference solutions, useful comparisons can still be drawn between simulations and the expected qualitative behaviour of the instability, with a motivation of reducing the clear suppression of the KHI observed when using traditional SPH. We additionally use equal mass particles across the simulation, making this setup particularly challenging for SPH schemes, but more applicable to most science applications. In §2.3.3 we present KHI simulations with a larger density

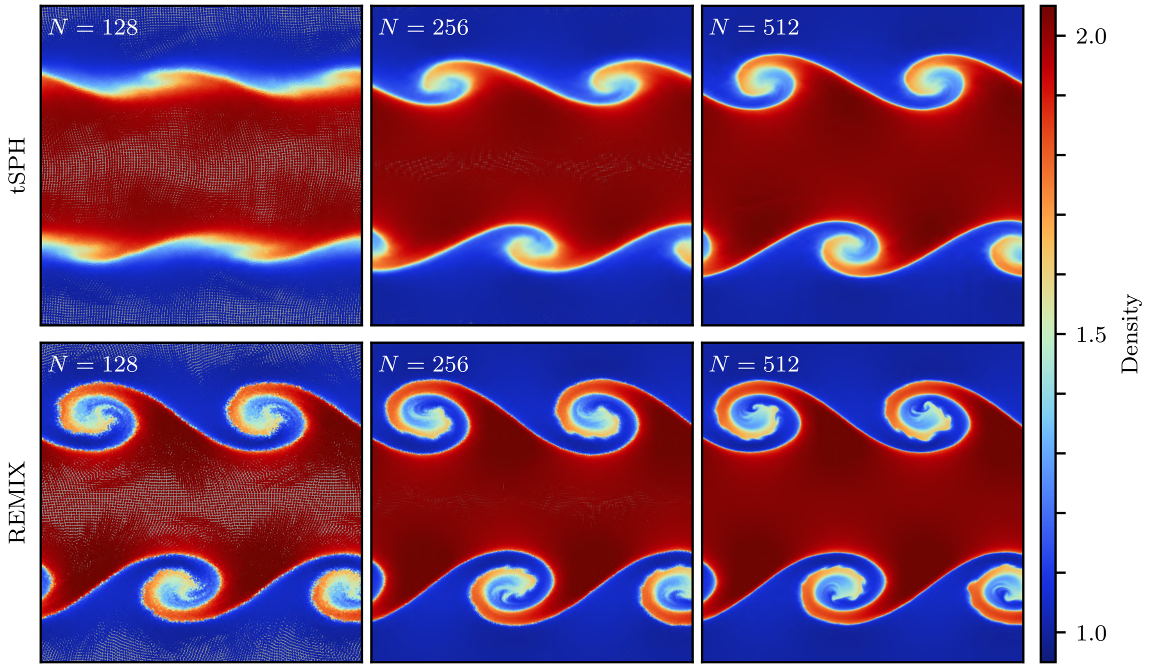

Figure 2.3: Ideal gas Kelvin–Helmholtz instabilities with smoothed initial density and velocity profiles. Columns correspond to simulations of different resolutions, with the top row showing results from simulations using tSPH and the bottom row from simulations using REMIX. We plot snapshots at $t = 2\ \tau_{\rm KH}$ from 3D, ideal gas KHI simulations. The density ratio between the two regions is 1:2. Here particles of different mass are used to match consistent initial particle spacing and volume across the simulation. Individual particles are plotted on a grey background and coloured by their density. Particles at all z are plotted, so the grey background is visible in regions that have maintained their grid alignment in z from the initial conditions.

ratio, a discontinuous interface, and equal mass particles. This system is even more challenging again for SPH schemes: both smoothing and discretisation errors are increased here due to the larger density-smoothing effects and the even more extreme local anisotropy in particle distribution at the interface. After considering these ideal gas scenarios, we present KHI simulations at interfaces between dissimilar, stiff materials in §2.3.4.

KHI with smooth initial conditions

McNally et al. (2012) present converged, high-resolution simulations of the early linear growth of the KHI. They use initial conditions with smooth initial velocity and density fields across the shearing interface. Similarly to Rosswog (2020), here

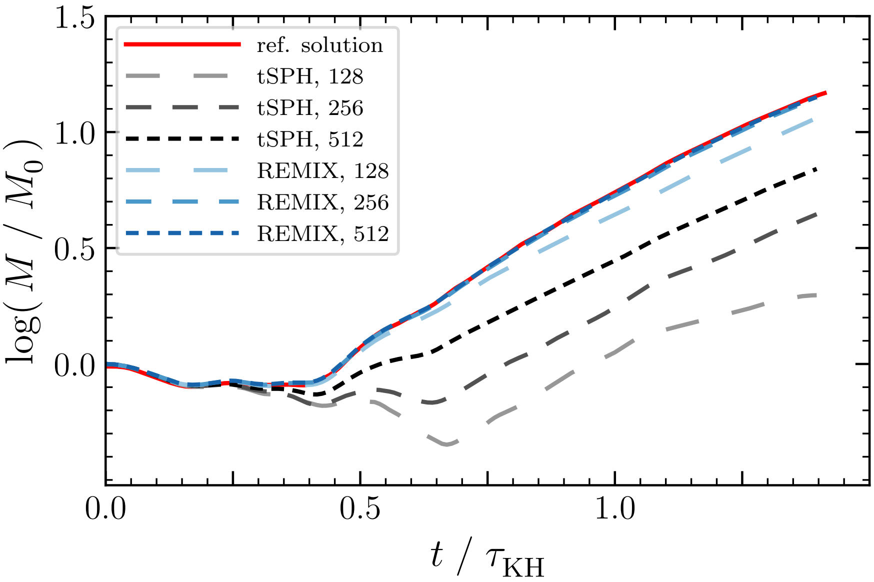

Figure 2.4: Evolution of mode amplitude, M, in Kelvin–Helmholtz simulations with smoothed initial density and velocity profiles. We plot mode growth for simulations with three different resolutions ($N = 128$, 256, 512) using both tSPH (grey, dashed) and REMIX (blue, dashed). Mode amplitude is normalised to the initial amplitude of the excited mode, M_0, and time is normalised to the characteristic timescale of KHI growth, $\tau_{\rm KH}$. The reference solution (red, solid) corresponds to the 4096 × 4096 simulation of McNally et al. (2012) using the Pencil code.

we use these smooth initial conditions, adapted to 3D, and use the mode growth of the reference solution of McNally et al. (2012) to quantitatively assess the accuracy of our numerical methods.

Particles are initialised in a 3D cubic lattice in a periodic box with $N \times N \times 18$ particles in x, y, z directions (i.e. a thin slice in the z direction relative to x and y). We run simulations with resolutions $N = 128$, 256, 512. Spatial dimensions are normalised to the size of the simulation box length in the x and y directions. A low density region of $\rho_1 = 1$ shears against a high density region of $\rho_2 = 2$ with speeds in the x direction of $v_1 = -0.5$ and $v_2 = 0.5$ such that the relative velocity is $|v_1 - v_2| = 1$. The regions are layered in y and have relative velocities in x. However, both density and shearing velocity are smoothed at the shearing interface such that initial densities are given by

$$\rho(y) = \begin{cases} \rho_1 - \rho_m e^{(y-0.25)/\Delta} & \text{for } 0.00 \leq y < 0.25\,, \\ \rho_2 + \rho_m e^{(0.25-y)/\Delta} & \text{for } 0.25 \leq y < 0.50\,, \\ \rho_2 + \rho_m e^{(y-0.75)/\Delta} & \text{for } 0.50 \leq y < 0.75\,, \\ \rho_1 - \rho_m e^{(0.75-y)/\Delta} & \text{for } 0.75 \leq y < 1.00\,, \end{cases} \tag{2.41}$$

and initial velocities in the x direction are given by

$$v_x(y) = \begin{cases} v_1 - v_m e^{(y-0.25)/\Delta} & \text{for } 0.00 \leq y < 0.25\,, \\ v_2 + v_m e^{(0.25-y)/\Delta} & \text{for } 0.25 \leq y < 0.50\,, \\ v_2 + v_m e^{(y-0.75)/\Delta} & \text{for } 0.50 \leq y < 0.75\,, \\ v_1 - v_m e^{(0.75-y)/\Delta} & \text{for } 0.75 \leq y < 1.00\,. \end{cases} \tag{2.42}$$

Here $\rho_m = (\rho_1 - \rho_2)/2$, $v_m = (v_1 - v_2)/2$, and $\Delta = 0.025$. Since particle positions are initialised in a single cubic lattice, particle masses are set by $m_i = \rho(y_i)/N^3$. Particle internal energies are set to give a pressure of $P(\rho, u) = 2.5$ across the simulation for an ideal gas with $\gamma = 5/3$. A small velocity perturbation, $v_y = 0.01 \sin(2\pi x/\lambda)$, is added in the y direction with wavelength $\lambda = 0.5$, to seed the primary instability.

The simulated KHI with these initial conditions is shown in Fig. 2.3. We plot particle densities at particle positions for simulations of resolution $N = 128$, 256, 512. Top row plots correspond to tSPH and bottom to REMIX. All snapshots are shown at simulation time $t = 2\ \tau_{\rm KH}$. Traditional SPH struggles to capture this instability at low resolutions. In REMIX simulations the seeded mode is not suppressed and grows at a close to resolution-independent rate. We find, however, that at later times secondary modes will eventually grow and disturb the evolution of the primary mode, so we do not observe strict convergence over long timescales. For an SPH scheme aiming to model an inviscid fluid with realistic turbulence-driven mixing, a compromise on this is difficult to avoid.

The evolution of the amplitude of the seeded mode is shown in Fig. 2.4, for these simulations. This quantity, M, is calculated from Eqns. 10–13 of McNally et al. (2012). We normalise the mode amplitude to $M_0 \equiv M(t = 0)$ to allow for more direct comparisons between simulations with different initial conditions, presented later. The reference solution is from the high-resolution 4096^2 cell KHI simulation performed by McNally et al. (2012) using the Eulerian mesh, finite-difference code Pencil. The mode growth of the tSPH simulations fall significantly short of the reference solution. This result is consistent with the SPH simulations used for comparisons by McNally et al. (2012). In contrast, the mid- and high-resolution REMIX simulations closely match the reference solution, and even the lowest resolution REMIX simulation is considerably closer to the reference solution than the highest resolution tSPH simulation.

KHI with discontinuous initial conditions

Next we consider KHI growth at an interface that is discontinuous in density and velocity. A shearing discontinuous interface is unstable to modes of all wavelengths, so noise- or error-seeded secondary modes will inevitably lead to a turbulent system and preclude numerical convergence. Although no converged reference solution exists for this problem, we use this system to qualitatively demonstrate the suppression in tSPH of both instability growth at, and mixing across, density discontinuities, and the effectiveness of REMIX in alleviating these issues. We deliberately constrain our analysis to low-resolution simulations, where the primary, intentionally seeded mode remains relatively undisturbed by secondary modes during the early growth of the instability (as discussed further in Appendix G).

Similarly to in §2.3.3, we consider shearing between a low density region of $\rho_1 = 1$ and a high density region of $\rho_2 \approx 2$ with relative speeds of $v_1 = -0.5$ and $v_2 = 0.5$. Here we initialise particles with a sharp discontinuity in both density and shearing velocity. The low density region is set up in the same cubic lattice as in the smoothed simulations of the previous section. We use particles of equal mass across the simulation. We refer to the resolution of these simulations by the effective resolution of the low density region: $N_1 = 128$, 256, 512; if the box were filled with a cubic lattice of the low density material only, then this lattice would consist of $N_1 \times N_1 \times 18$ particles in x, y, z directions. Particles in the high density region are arranged in a cubic lattice of shorter spacing. A cubic lattice corresponding to a density $\rho_2 = 2$ is adjusted to allow a continuous grid in the x dimension of

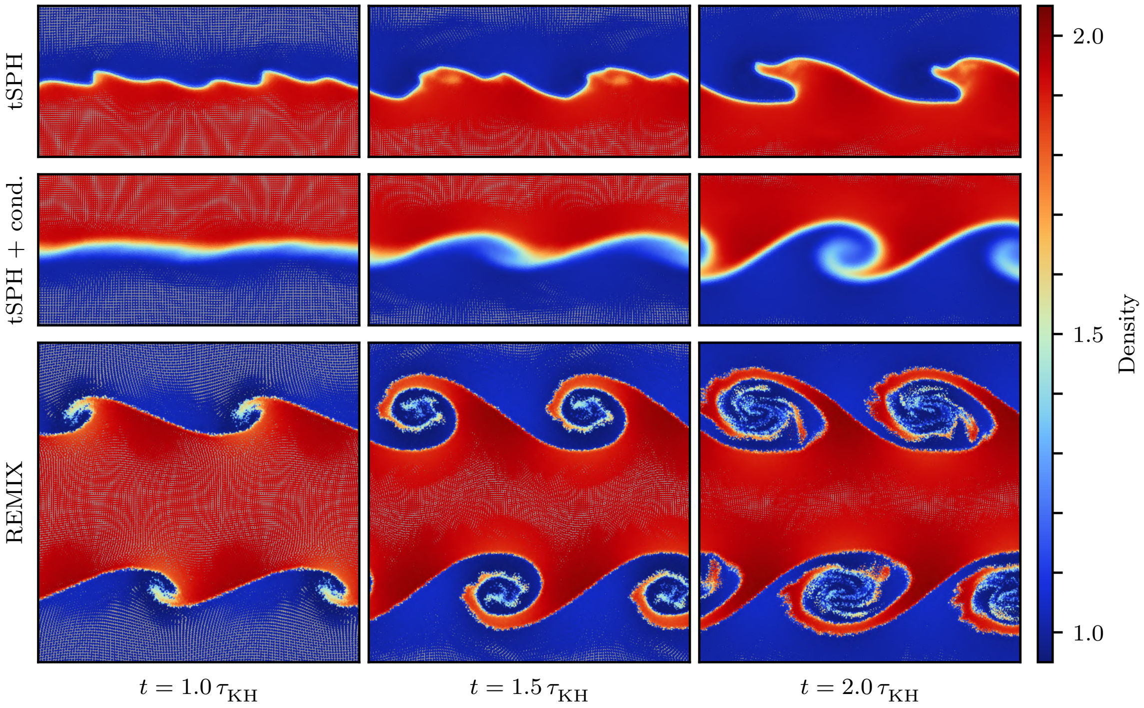

Figure 2.5: Growth of 3D Kelvin–Helmholtz instabilities in the more challenging case of discontinuous initial density and velocity profiles and equal mass particles. Columns show snapshots at different times, with the top rows showing results from simulations using tSPH – without and with artificial conduction – and the bottom row from simulations using REMIX. These simulations are both relatively low resolution, with $N_1 = 128$. The density ratio between the two regions is 1:1.91.

the periodic box[18]. The spacing of particles in z is slightly adjusted away from a perfectly cubic lattice such that particle spacing in this dimension is also continuous across the boundary of the box. The regions are shifted in the y direction such that the layers of particles across the interface from each other, directly adjacent to, and parallel with, the discontinuity are separated by the mean of the two grid-spacings. The size of the simulation box is adjusted in the y direction to compensate for this and to maintain continuity across boundaries of the periodic box. The density ρ_2 is recalculated based on these grid modifications and the use of equal mass particles. To satisfy these conditions, in the high density region we use $\rho_2 = 1.91$ in a lattice

---

[18] We also enforce that the effective resolution in this region, N_2, is divisible by 4 (the number of vortices formed by the evolution of the seeded mode) to avoid the possibility of asymmetric evolution of individual vortices triggering an early onset of secondary modes. In practice this has no noticeable effect here, but similar considerations do matter for the Rayleigh–Taylor instability simulations we examine in §2.3.5 and §2.3.6.

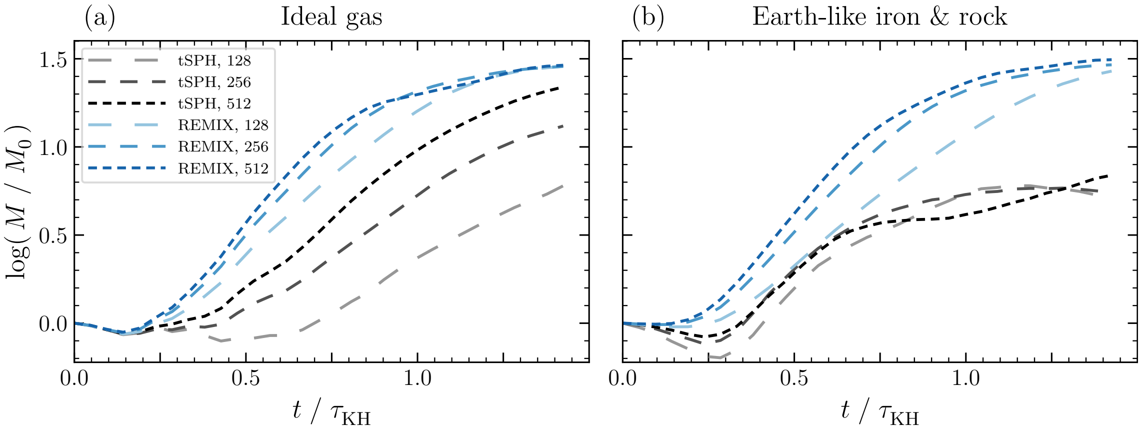

Figure 2.6: Evolution of Kelvin–Helmholtz instability mode amplitude, M, in sharp-discontinuity, equal mass KHI simulations with (a) an ideal gas EoS (§2.3.3) and (b) between dissimilar, stiff EoS (§2.3.4). Mode growth in simulations using both tSPH (grey, dashed) and REMIX (blue, dashed) are shown for three different resolutions (N_1 = 128, 256, 512). The mode amplitude is normalised to the initial amplitude of the excited mode, M_0, and time is normalised to the characteristic timescale of KHI growth, $\tau_{\rm KH}$, which are both different for each case.

with, for example, $N_2 = 160$ and 22 particles in the z direction for $N_1 = 128$. Initial internal energies are calculated such that particles have a uniform initial pressure[19] of $P(\rho, u) = 2.5$ by the ideal gas equation with $\gamma = 5/3$. We seed a small velocity perturbation, $v_y = 0.01 \sin(2\pi x/\lambda)$, in the y direction with $\lambda = 0.5$.

In Fig. 2.5 we show the growth of these KHIs simulated using tSPH, tSPH with artificial conduction, and REMIX. We plot individual particles, coloured by their densities, at three times through the evolution of the instability. In the tSPH simulations, surface tension-like effects act both to suppress the growth of the instability and to prevent mixing of particles across the interface. As noted by Agertz et al. (2007), particles form ordered bands with large gaps at the interface, which act as barriers to mixing. Artificial conduction helps to enable some mixing on the particle scale, allowing the boundary to become diffuse with time. However, the evolution of the instability is slow, as can be seen when comparing with the similar sharp-interface KHI simulations of Hopkins (2015), performed with their improved methods, at comparable scaled times (their Fig. 21). While we note that differences

---

[19] We note here that the density used in these initial conditions are unsmoothed, so the tSPH simulations will not be in pressure equilibrium due to their smoothing of the densities at the discontinuities.

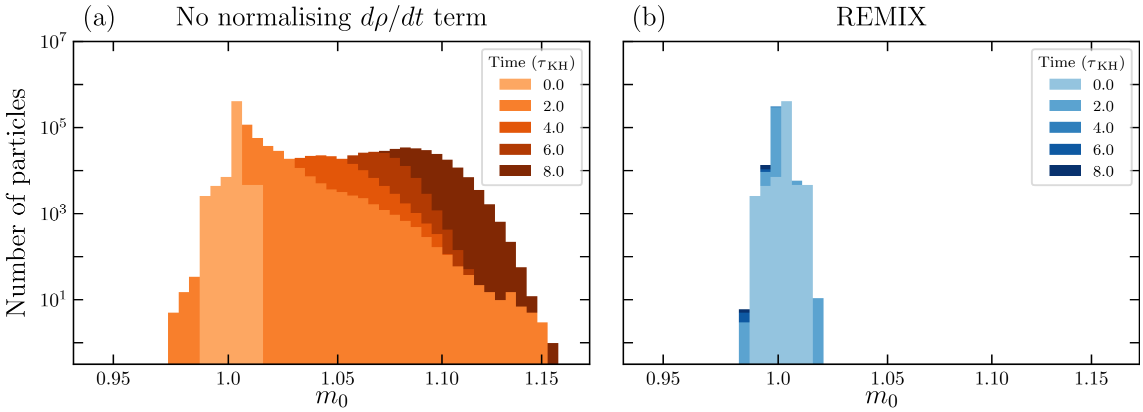

Figure 2.7: The evolution of the distribution of the m_0 kernel geometric moment in ideal gas KHI simulations with sharp discontinuities. $m_0 \approx 1$ corresponds to a local distribution of densities that reflects the particle configuration. Plots show results from simulations with resolution $N_1 = 128$ using the REMIX scheme both (a) without and (b) with the normalising term in the density evolution.

in initial conditions construction means that we cannot make direct comparisons, the growth of the instability in both traditional cases is clearly too slow. The REMIX simulation shows a clear improvement: not only do the characteristic vortices of the KHI form without impedance by surface tension-like effects, but interfaces are maintained as sharp discontinuities as the system evolves. Particles do not align themselves in bands separated by gaps that would prevent mixing across the discontinuity.

The mode amplitude growth of these KHIs and equivalent higher resolution simulations are plotted in Fig. 2.6(a). Since this system is constructed differently from that in §2.3.3, we cannot draw direct comparisons between these results and the converged reference solution for a smoothed interface. For example, the instability grows more quickly in this case where the shearing velocity is discontinuous. However, we do observe qualitatively similar behaviour when comparing Fig. 2.6(a) with Fig. 2.4. The seeded mode grows more quickly in REMIX simulations than in those using tSPH. The early growth rate of modes in REMIX simulations is slightly steeper as resolution is increased, mirroring the behaviour of the analogous simulations in Fig. 2.4. The approach of the mode evolution of tSPH simulations towards the REMIX simulations is also similar here, and again, the lowest resolution REMIX simulation grows more quickly than the highest resolution tSPH simulation. Despite this behaviour with increased resolution, high-resolution tSPH simulations still fail to form spiralling plumes, as surface tension-like effects continue to dominate, as shown in Appendix G.

The effect of the normalising term (§2.2.7) in a KHI simulation with sharp discontinuities is demonstrated in Fig. 2.7. Without it, as the simulation evolves, m_0 of some particles drifts away from 1, the value corresponding to normalisation of the unmodified kernel (see Eqn. 2.38). The normalising term ties the density evolution to kernel normalisation, so as the system evolves, volume elements continue to accurately build up the continuum over which the kernel function is normalised. In these simulations, the drift in m_0 does not noticeably affect the simulation outcome, however, in §2.3.9, we show an example where the inclusion of the normalising term is necessary to simulate a system in hydrostatic equilibrium.

KHI with a large density ratio

Capturing the KHI at interfaces in fluids with a large density jump is additionally challenging for SPH. More smoothing in the density estimate and larger discretisation error, in equal mass particle simulations, will make surface tension-like effects stronger at larger density contrasts. Additionally, using artificial conduction as a method for correcting density discontinuity is not as effective at larger jumps in density (Price, 2008). Our initial conditions aim to follow those of Price (2008) with a density ratio of 1:10, however, we continue to use 3D simulations to validate our methods for more typical applications.

Here we construct initial conditions similarly to §2.3.3: sharply discontinuous in both density and shearing velocity. The low density region is constructed exactly equivalently with $\rho_1 = 1$, while resolution is increased in the high density region, following the same method as outlined previously, such that this region has a density of $\rho_2 = 10.4$. Speeds in the x direction are again set to $v_1 = -0.5$ and $v_2 = 0.5$, however the wavelength of the initial perturbation in the y direction is decreased to $\lambda = 0.128$, although with the same amplitude of 0.01 (Price, 2008). Comparisons of resolution can not be directly drawn to the previous section, as here fewer particles will make up individual vortices at a given time due to the decreased perturbation

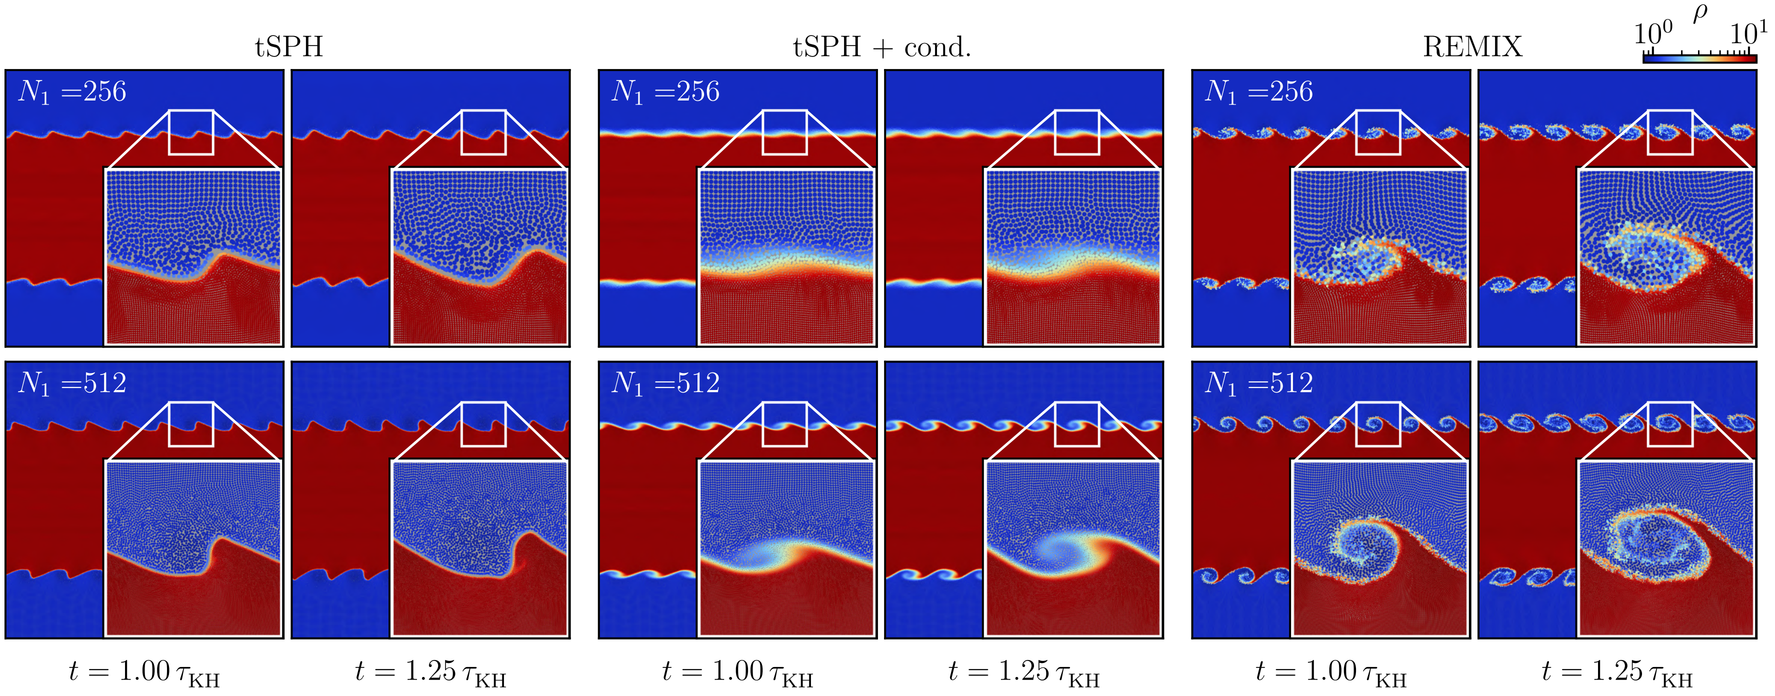

Figure 2.8: Growth of Kelvin–Helmholtz instabilities between layers with a large density contrast. We plot snapshots from 3D, ideal gas KHI simulations with discontinuous initial density and velocity profiles. We show results from simulations carried out using tSPH, tSPH with artificial conduction, and REMIX. Snapshots are plotted at two times, from simulations with two resolutions. The density ratio between the two regions is 1:10.4 and particles have the same mass across the simulation. Insets show a magnified view of a KHI plume.

wavelength.

In Fig. 2.8, we plot snapshots showing the evolution of these initial conditions in tSPH, tSPH with conduction, and REMIX simulations for two resolutions. Due to the lower wavelength of the seeded mode compared with that in previous sections, we consider simulations with overall higher resolutions, although this does not necessarily correspond to higher resolution in each individual vortex, which now occupies a smaller region in the simulation box. The instability fails to grow with tSPH and grows only slowly in the higher resolution simulation with conduction. However, the instability is captured successfully with REMIX, in particular at the higher-resolution, where we capture spiralling within the plume.

2.3.4 Kelvin–Helmholtz instability – Earth-like iron & rock

Since the evolution of the KHI is predominantly inertial, we expect instabilities to grow similarly between shearing fluids of different materials, represented in our simulations as inviscid fluids only differing in the calculation of pressures and sound speeds through the EoS (§2.1.1). We construct similar initial conditions to those used in §2.3.3, but using the ANEOS $Fe_{85}Si_{15}$ (iron) and forsterite (rock) EoS with densities and pressures comparable with those of the Earth's core-mantle interface (Stewart et al., 2020).

We simulate the KHI at a discontinuity between low-density rock at $\rho_1 = 5000$ kg m^{-3} and high-density iron at $\rho_2 = 9550$ kg m^{-3}. Particles are placed in a periodic box in a configuration exactly matching that of §2.3.3. These simulations use particles of equal mass. Spatial dimensions are scaled such that the box spans a length of $1\ R_\oplus = 6371$ km in the x and y dimensions. The velocities in x of the two layers are initialised to $v_1 = -10^{-4}\ R_\oplus\ \mathrm{s}^{-1}$, $v_2 = 10^{-4}\ R_\oplus\ \mathrm{s}^{-1}$ and the seeded mode has the form $v_y = 0.01|v_1 - v_2|\sin\left(2\pi x/\lambda\right)$ with $\lambda = 0.5\ R_\oplus$. Initial internal energies are calculated through each material's EoS such that the regions are in pressure equilibrium with $P(\rho, u) = 1.2 \times 10^{11}$ Pa.

In Fig. 2.9 we show the evolution of a KHI with these initial conditions using tSPH and REMIX. In the tSPH simulation, surface tension-like effects are strong. Undesired smoothing of the discontinuity in the SPH density estimate combined with

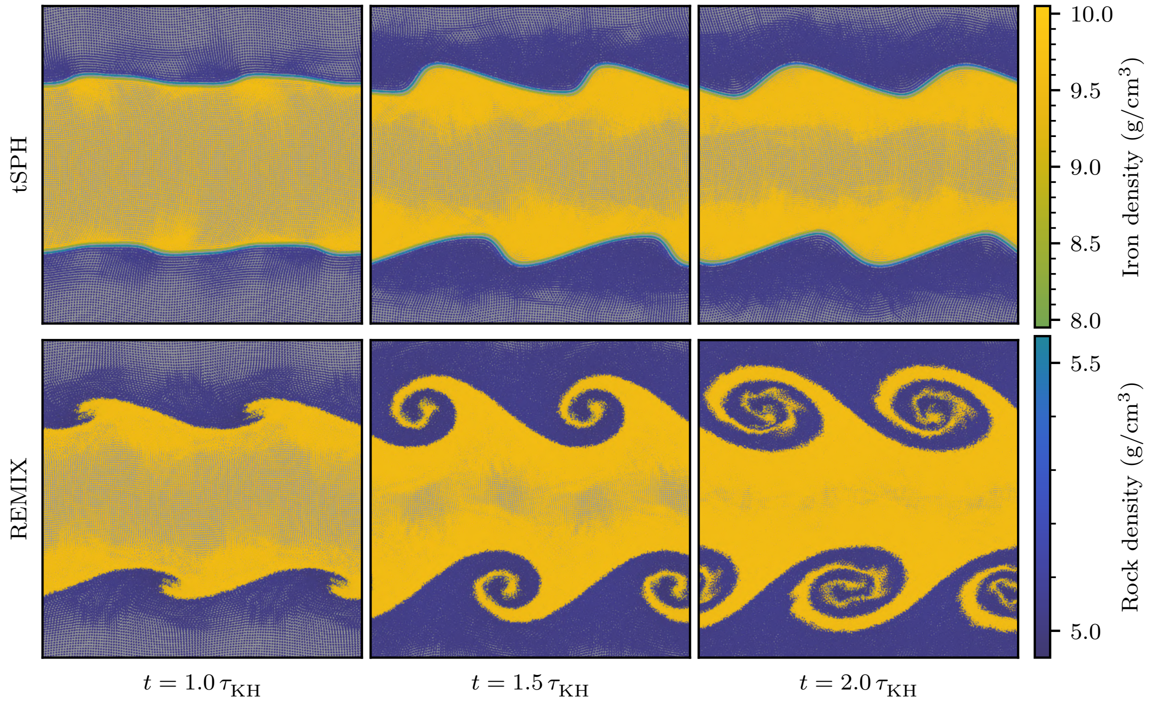

Figure 2.9: Kelvin–Helmholtz instability growth between dissimilar, stiff materials. We plot snapshots from 3D KHI simulations with multiple, complex equations of state at densities and pressures representing those at material boundaries within the Earth. Columns show snapshots at different times with the top row showing results from simulations using tSPH and the bottom row using REMIX. The initial density and velocity profiles are discontinuous and particles have equal mass. These simulations are both relatively low resolution, with $N_1 = 128$. Individual particles are plotted on a grey background and coloured by their material type and density. Particles at all z are plotted, so the grey background is visible in regions that have maintained their grid alignment in z from the initial conditions.

the stiff equations of state leads to strong artificial forces at the interface, which both prevent mixing of particles of different materials and strongly suppress the growth of the instability. These effects as well as their contributions from zeroth-order error in the momentum equation are addressed in the construction of the REMIX SPH scheme, so the instability is allowed to grow and particles of different materials are able to intermix in a qualitatively similar way to the ideal gas cases.

The mode amplitude growth of these simulations is plotted in Fig. 2.6(b). We find strong quantitative similarities between these and the mode growth of the ideal gas simulations potted in Fig. 2.6(a). Although we have no experimental or analytical predictions for the growth of the KHI in these conditions and with these materials, we find that: (1) spurious surface tension analogous to that in tSPH KHI simulations

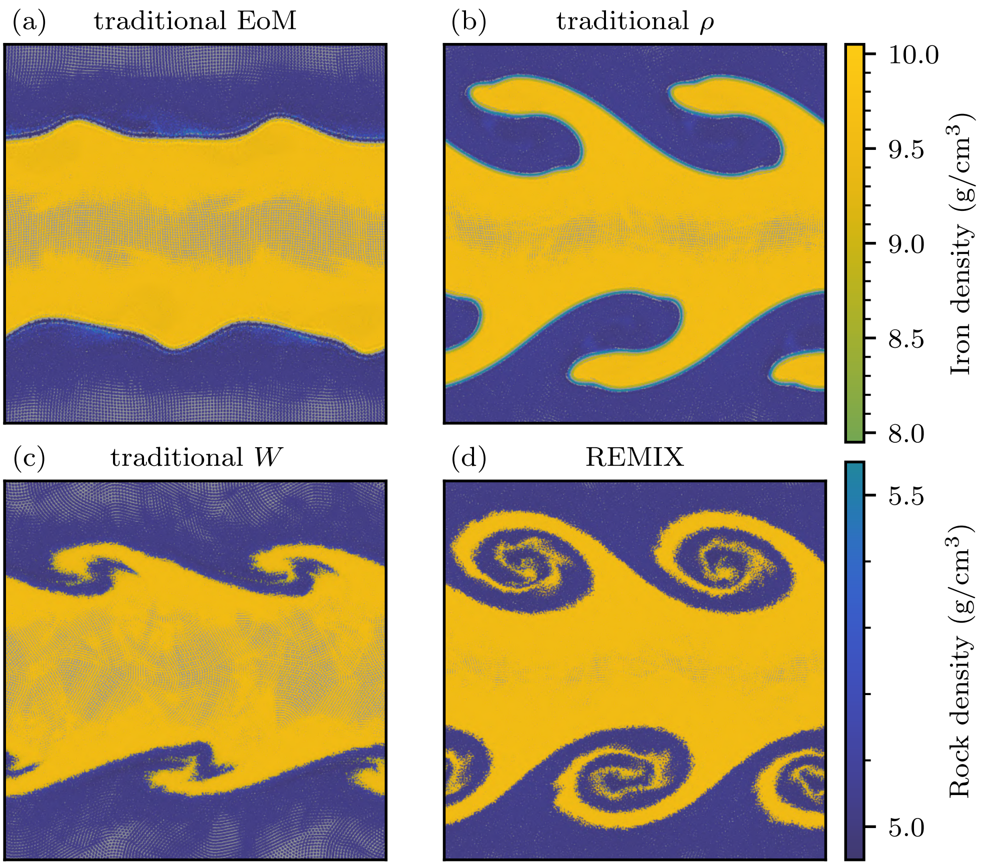

Figure 2.10: Kelvin–Helmholtz instability simulations illustrating the interplay between multiple component methods of the REMIX scheme. Here we show results from REMIX simulations when using: (a) a more traditional form of the equations of motion; (b) the traditional, integral-form of the density estimate, instead of the evolved density; (c) an unmodified Wendland C^2 kernel, instead of the linear-order reproducing kernels; (d) the full REMIX SPH scheme. Removing any one of these affects the growth of the instability significantly. We plot snapshots at $t = 2\ \tau_{\rm KH}$ from 3D simulations with multiple, stiff equations of state at densities and pressures representative of those at the core-mantle boundary within the Earth.

with ideal gas is also clearly visible and strong in tSPH simulations with multiple materials; (2) the construction of the REMIX scheme is general in, and shown to be effective in, its reduction of established sources of error in the SPH formalism; (3) without any tuning of the method to material-specific boundaries, improvements that alleviate surface tension-like effects in ideal gas KHI simulations also allow the KHI to form in a qualitatively similar manner in the multi-material case.

To achieve these improved results of the REMIX scheme demonstrated in Fig. 2.9, we require interplay between a combination of its constitutive methods (§2.2). We use this KHI with iron & rock to highlight the importance of individual methods included in the REMIX SPH scheme as, while their effects are visible in all simulations, they present particularly clearly in this case. Fig. 2.10 shows Earth-like KHI simulations that use the REMIX SPH scheme with different ones of its constituent methods removed from the construction and reverted to its traditional SPH analogue in each panel. We show simulations that: (a) use a more standard form of the equations of motion with equal-valued free functions (§2.2.2); (b) use the integral rather than differential form of the density estimate (§2.2.1); (c) use an unmodified Wendland C^2 kernel rather than linear-order reproducing kernels (§2.2.3); (d) the full REMIX SPH scheme. Taking a more traditional approach in any one of these methods leads to much stronger surface tension-like effects, such that only the full scheme enables the expected spirals to form. The improvements of the REMIX scheme are in many cases due to interplay between its constitutive methods all together, rather than individual components solving separate issues.

2.3.5 Rayleigh–Taylor instability – ideal gas

We next consider the Rayleigh–Taylor instability (RTI) as an additional scenario to test the treatment of instability growth and mixing, which, unlike the previous tests, also includes gravity.

The RTI arises at the interface between a high density fluid being displaced by a low density fluid (Chandrasekhar, 1961). We simulate the gravity-driven growth of the RTI, where a layer of high density fluid is initially positioned above a layer of low density fluid (relative to the downward direction of gravity). Hydrostatic equilibrium is disturbed by a small velocity perturbation. Similarly to the KHI, surface tension-like effects in traditional SPH formulations strongly suppress the growth of the RTI.

Our initial conditions are based on those of Frontiere et al. (2017). However, as with the KHI tests, these simulations are carried out in 3D, with particles of equal mass, and without deliberate smoothing of the initial density discontinuity. Particles are placed in a periodic simulation domain in two cubic lattices, each a square in the x, y dimensions and thin in z. The box has dimensions of 0.5, 1 in the x and y directions, with a thin and resolution-dependent z box size. The low density

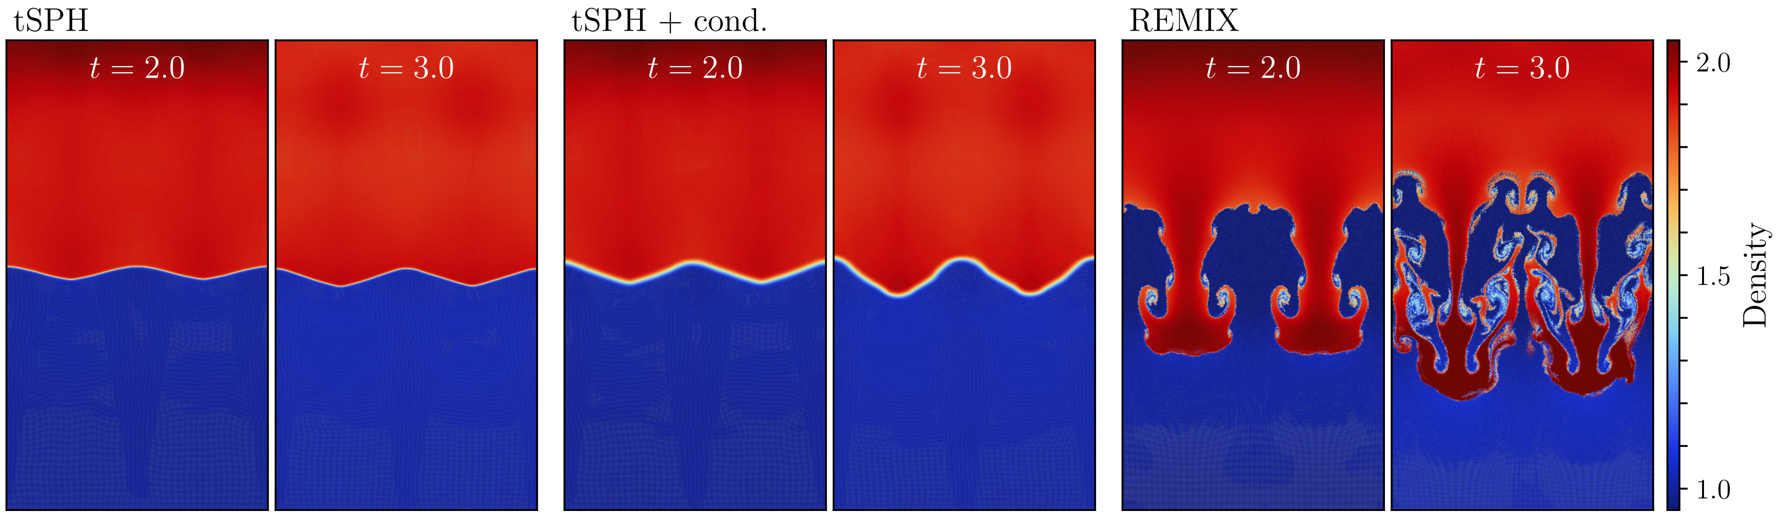

Figure 2.11: Rayleigh–Taylor instabilities in an ideal gas. The RTIs are shown at two times, in simulations using tSPH, tSPH with artificial condition, and REMIX. These simulations have a resolution of $N_1 = 256$, equal mass particles, and are performed in 3D. The regions of fixed particles at the top and bottom of the simulations have been cropped from the figure; their positions and densities do not change.

region has $N_1 \times N_1 \times 18$ particles with density $\rho_1 = 1$ and occupies the bottom half of the domain. The high density region is constructed similarly to that in §2.3.3, giving a density of $\rho_2 = 1.91$ for the upper region while also ensuring a lattice that is consistent across the periodic simulation box edges. Particles in the top and bottom 0.05 of the box are fixed in place throughout the course of the simulation. Initial internal energies are set to satisfy hydrostatic equilibrium using an ideal gas EoS with $\gamma = 7/5$, constant gravitational acceleration $g = -0.5$, and a pressure at the interface of $P_0 = \rho_2/\gamma$. Particles are initially at rest, other than an initial velocity perturbation that seeds the instability,

$$v_y(x,y) = \begin{cases} \delta_y \left[1 + \cos\left(8\pi\left(x + 0.25\right)\right)\right]\left[1 + \cos\left(5\pi\left(y - 0.5\right)\right)\right] & \text{for } 0.3 < y < 0.7 \, , \\ 0 & \text{otherwise.} \end{cases} \tag{2.43}$$

We use a perturbation amplitude of $\delta_y = 0.025$.

In Fig. 2.11, we show snapshots from RTI simulations with resolution $N_1 = 256$, simulated using tSPH, tSPH with artificial conduction, and REMIX. The growth of this instability is strongly suppressed, even with artificial conduction. REMIX is able to capture the growth of the RTI well. Additionally, we are able to maintain discontinuities as the simulation evolves. As in the KHI, these discontinuities are inherently unstable to perturbation modes of all wavelengths and so we see the growth of secondary, unseeded KHIs and RTIs that contribute to an onset of turbulent mixing. As the simulation progress, we observe turbulence driving mixing on the particle scale, the scale of the primary instability, and in between.

2.3.6 Rayleigh–Taylor instability – Earth-like iron & rock

We now consider the treatment of the RTI at an interface between different materials. The stiff iron & rock EoS that we use makes this an even more challenging scenario for traditional SPH.

The high density iron layer is placed above the low density rock layer. The particle configuration is constructed just as in the ideal gas case above, since the density

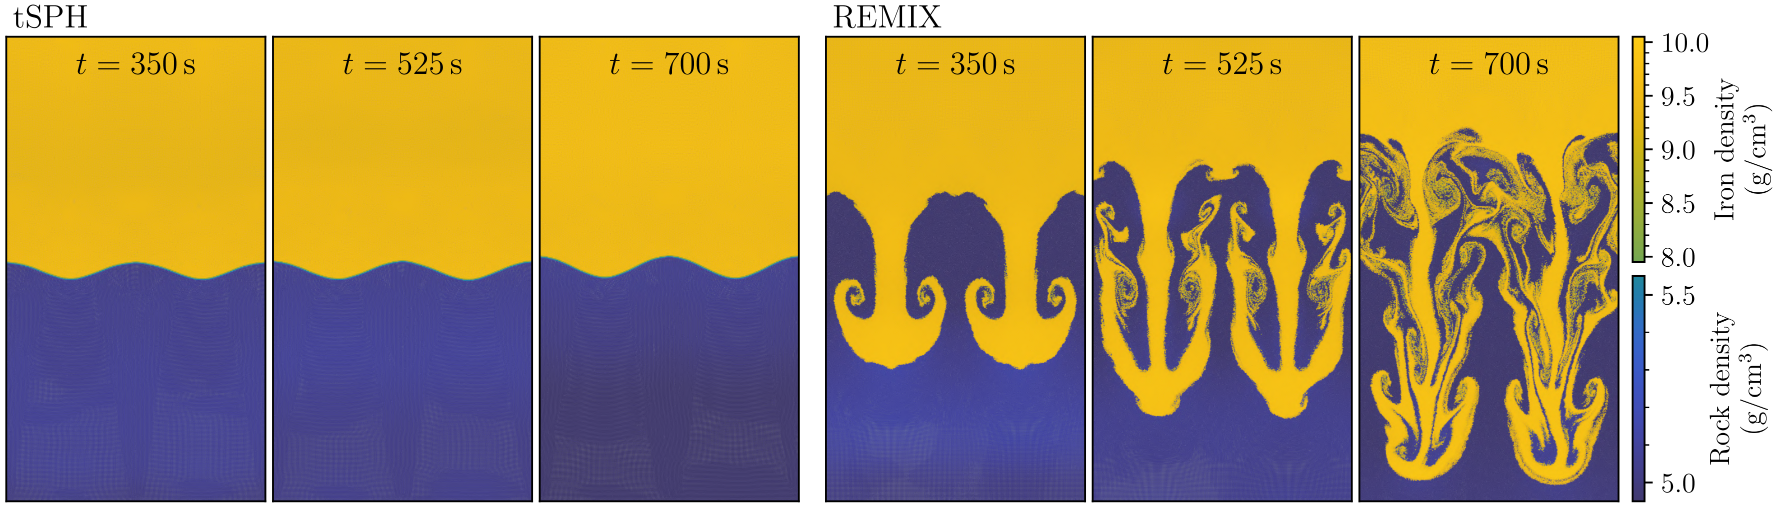

Figure 2.12: Rayleigh–Taylor instabilities between dissimilar, stiff materials. The RTIs are shown at three times, in simulations using tSPH and REMIX. These simulations have a resolution of $N_1 = 256$, equal mass particles, and are performed in 3D.

ratio is taken to be the same. However the box is scaled to have dimensions 0.05 $R_\oplus$ and 0.1 $R_\oplus$ in the x and y dimensions. The velocity perturbation is similar, although scaled to the box size and with an amplitude $\delta_y = 2.5 \times 10^{-5}$ $R_\oplus$ s^{-1}. Again, particles are initially in hydrostatic equilibrium, other than due to the seeded perturbation. Internal energies are chosen to satisfy this for the constant gravitational acceleration $g = -9.9$ m s^{-2} and interface pressure $P_0 = 120$ GPa, representative of the gravitational acceleration and pressure at the Earth's core-mantle boundary.

In Fig. 2.12 we show snapshots from RTI simulations with Earth-like materials with resolution $N_1 = 256$, with tSPH and REMIX. The RTI does not grow in the tSPH simulation. In contrast, the behaviour of the REMIX simulation is similar to the equivalent ideal gas case: unimpeded evolution of the instability, mixing at different length scales, onset of turbulence, and growth of unseeded secondary modes.

2.3.7 Blob test

In a physical system, mixing due to fluid instabilities is typically much less controlled and isolated than in the deliberately seeded scenarios of the previous sections. The "blob test" (Agertz et al., 2007) is used to investigate the treatment of turbulent mixing at density discontinuities due to unseeded instabilities.

A spherical cloud of high-density fluid, initially at rest, is placed in an uniform flow of low-density fluid. Emergent Kelvin–Helmholtz and Rayleigh–Taylor instabilities at the interface, as well as ram-pressure stripping, should act to break up the cloud, driving its evolution to a well-mixed state. As with the instability tests presented in previous sections, traditional SPH schemes struggle to capture instability growth at density discontinuities and so the mixing of the cloud into the surrounding fluid is strongly suppressed. Typically, blob tests are carried out in a supersonic regime, where interactions between shock waves and the cloud can also be assessed, applicable to a range of astrophysical scenarios. However, here we additionally simulate blob tests in a subsonic regime to demonstrate the ability of the REMIX SPH scheme in capturing subsonic turbulent mixing, which is even more strongly suppressed in the tSPH formalism (Bauer & Springel, 2012).

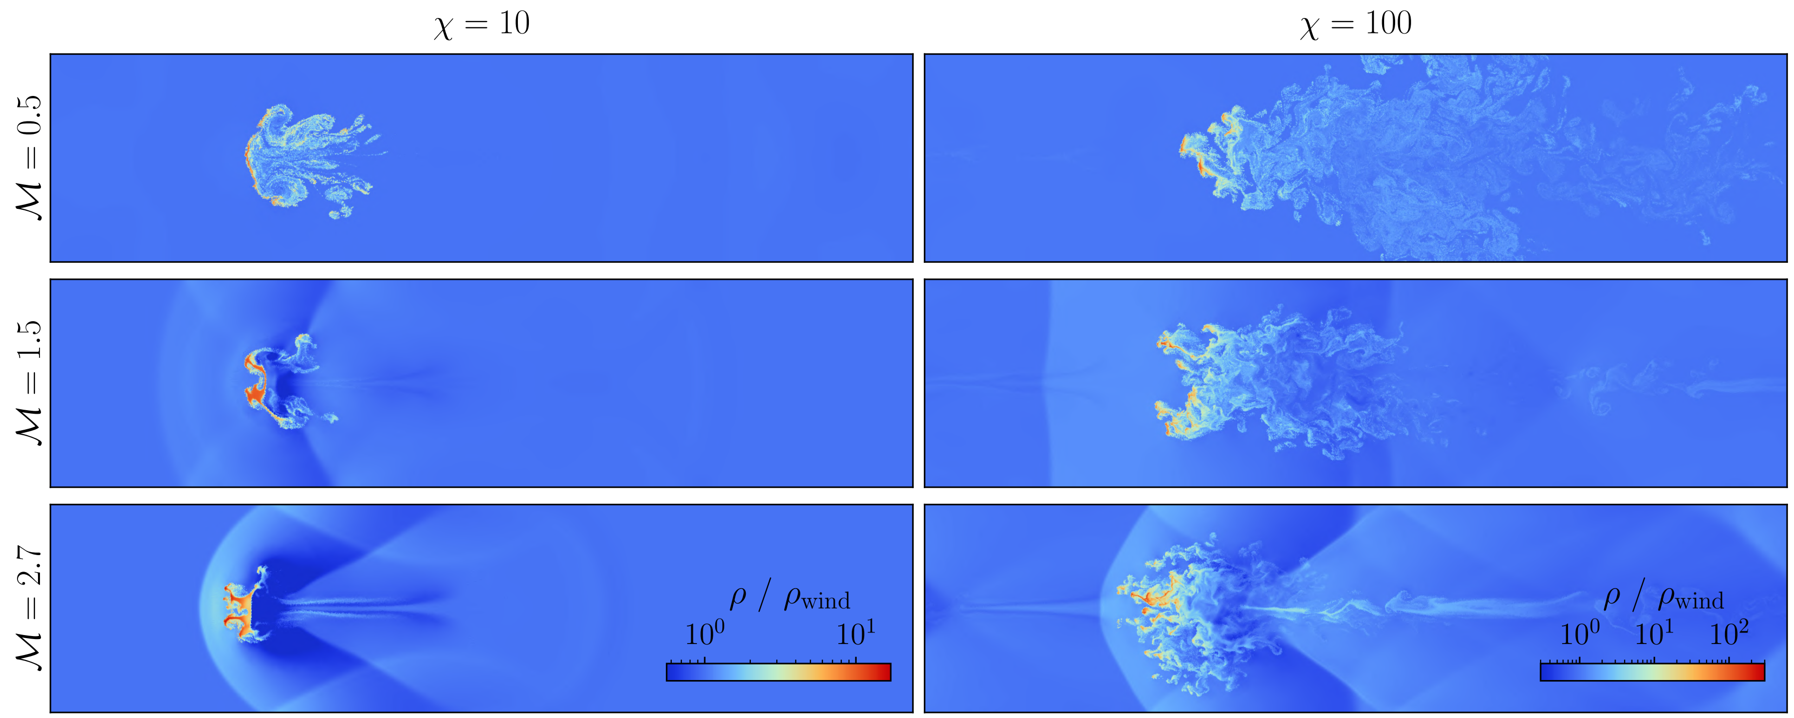

Figure 2.13: REMIX simulations of a high-resolution ($N = 256$) 3D blob test, at time $t \approx 5\ t_{\rm cc}$. The disruption of the cloud is shown in scenarios with two different initial density contrasts (columns) and three Mach numbers (rows). Individual particles in a central cross-section in z are plotted, and coloured by their density, ρ, relative to the wind density, $\rho_{\rm wind}$.

Braspenning et al. (2023) compare blob test simulations using seven hydrodynamical solvers, including SPH schemes and mesh-based methods. We reproduce their initial conditions to allow direct comparisons with their simulations. Particles are placed in a 3D periodic box with axes aligned such that the initially uniform wind flows in the x direction. The length of the box in y and z is chosen to be 1 pc, and the length in the x direction is 4 pc. Particles in the ambient wind are initially placed in a cubic lattice with $4N \times N \times N$ in the x, y, z directions, where we parameterise the simulation resolution by N. We carry out simulations with $N = 16$, 32, 64, 128, 256. Particles in both the cloud and surrounding wind have equal masses and so particles in the cloud are placed in a cubic lattice of higher number density corresponding to the chosen density contrast. We simulate blob tests with initial density contrasts $\chi = 10$, 100 and the initial density of the surrounding medium is 10^{-4} $\mathrm{m_p}$ $\mathrm{cm^{-3}}$, where $\mathrm{m_p}$ is the proton mass. Clouds are spherical and have a radius of $R_{\mathrm{cloud}} = 0.1$ pc. Both the cloud and surrounding medium are an ideal gas with $\gamma = 5/3$ and internal energies are chosen so that the cloud and surrounding medium are in pressure equilibrium with each other and the cloud has an initial temperature of 10^4 K. We carry out simulations with three wind speeds, characterised by the Mach number $\mathcal{M} \equiv v_{\mathrm{wind}}/c_{\mathrm{wind}}$: $\mathcal{M} = 1.5$ for a direct comparison to the simulations of Braspenning et al. (2023), $\mathcal{M} = 2.7$ the value used most frequently in validating hydrodynamic methods (Agertz et al., 2007; Frontiere et al., 2017), and $\mathcal{M} = 0.5$ to test mixing in the subsonic regime. We use units of the cloud crushing timescale

$$t_{\mathrm{cc}} = \frac{\sqrt{\chi}\, R_{\mathrm{cloud}}}{v_{\mathrm{wind}}} \,, \tag{2.44}$$

to compare simulations with different initial density contrasts and wind speeds and for direct comparisons with the results of Braspenning et al. (2023).

In Fig. 2.13 we plot particle densities from a central cross-section of REMIX blob test simulations with $N = 256$. Results from simulations with three initial wind speeds and two initial density contrasts are shown for a time $t \approx 5\ t_{\mathrm{cc}}$. The middle row therefore corresponds directly to results from simulations plotted in Fig. 1 of Braspenning et al. (2023). REMIX captures disruption of the cloud in both

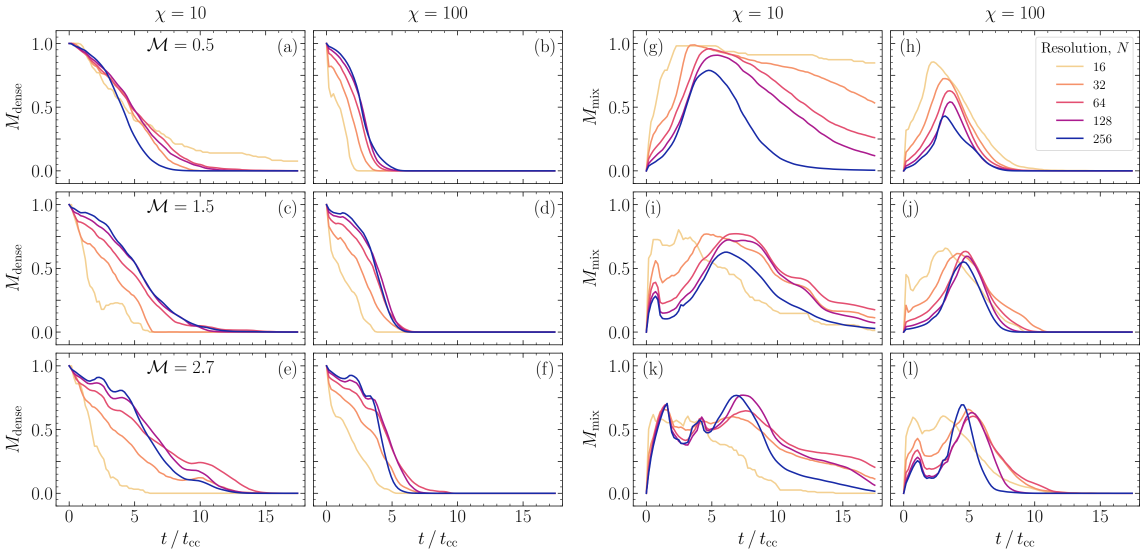

Figure 2.14: Time evolution of the mass of dense gas, $M_{\rm dense}$ (a–f), and the mass of intermediate-temperature gas, $M_{\rm mix}$ (g–l), in blob test simulations. These quantities are plotted for simulations for two different initial density contrasts (χ; columns) and three Mach numbers ($\mathcal{M}$; rows). Line colour corresponds to simulation resolution.

a subsonic and supersonic regime. With time, the cloud reaches a well-mixed state with the surrounding medium. This contrasts with the Braspenning et al. (2023) SPH simulations, in which clouds with $\chi = 100$ do not fully mix (their Fig. A1). Additionally, the onset of turbulence does not produce a highly symmetric structure like that provoked by the use of a regular grid in simulations using an adaptive mesh refinement (AMR) method (seen most clearly in Figs. 5 and 6 Braspenning et al. (2023)).

To facilitate direct quantitative comparisons to the simulations of Braspenning et al. (2023), we consider the evolution of the mass of dense gas and the mass of intermediate-temperature gas. The mass of dense gas, $M_{\rm dense}$, is defined as the sum of masses of particles with density above a threshold of $\rho_{\rm cloud}/\,3$, where $\rho_{\rm cloud}$ is the initial cloud density. The mass of intermediate-temperature gas, $M_{\rm mix}$, is defined as the sum of particle masses, m_i, of particles whose temperature, T_i, lies within half the logarithmic temperature range between the cold cloud and the hot wind, centred on the geometric mean temperature, i.e.

$$M_{\rm mix} = \sum_i m_{{\rm mix},\,i}\ , \tag{2.45}$$

with $m_{\mathrm{mix},\,i} = m_i$ for $\log(T_{\mathrm{mix}}) - \frac{1}{4}\log(\chi) < \log(T_i) < \log(T_{\mathrm{mix}}) + \frac{1}{4}\log(\chi)$ and 0 otherwise, where T_{mix} is the geometric mean of the cloud and wind temperatures, $T_{\mathrm{mix}} = \sqrt{T_{\mathrm{cloud}} T_{\mathrm{wind}}}$. We normalise both M_{dense} and M_{mix} to the initial cloud mass.

The evolution of M_{dense} and M_{mix} for simulations varying initial wind speed, initial density contrast, and resolution is plotted in Fig. 2.14. We find that REMIX is able to capture the disruption of the cloud in all these simulations, as shown in Fig. 2.14(a–f). The middle row corresponds directly to Figs. 2 and 3 of Braspenning et al. (2023). We see strong similarities between the behaviour of our REMIX simulations and the simulations of Braspenning et al. (2023) with hydrodynamic solvers that they find demonstrate good mixing. The evolution of these quantities is well parameterised by the cloud crushing timescale for this range of Mach numbers, with features appearing at approximately the same scaled time for all rows. Increasing resolution results in behaviour that indicates an approach towards numerical convergence for both M_{dense} and M_{mix}, despite the scenario itself being highly turbulent with no true converged solution.

2.3.8 Evrard collapse

The Evrard collapse (Evrard, 1988) considers the collapse of an isothermal, spherical cloud of gas under its self-gravity. A shock is formed and moves outwards as the cloud collapses. We use this test to investigate the coupling of gravity and hydrodynamics, with large transformations of energy between gravitational, kinetic and thermal forms.

Initial conditions are constructed similarly to Borrow et al. (2022). We place $\sim 10^7$ equal mass particles of ideal gas with $\gamma = 5/3$ and $u = 0.05$ in a spherical cloud of density profile $\rho(r) = 1/(2\pi r)$, where r is the radial distance from the cloud centre. The total cloud mass and radius are given by $M = 1$ and $R = 1$, and the gravitational constant is set to $G = 1$. Particle positions are chosen randomly, following Borrow et al. (2022), to satisfy the initial density profile. This method of choosing positions results in particles quickly readjusting to a glass-like structure, therefore experiencing divergences that lead to a seeding of noise in internal energies and densities. Both artificial diffusion and the normalising density evolution term

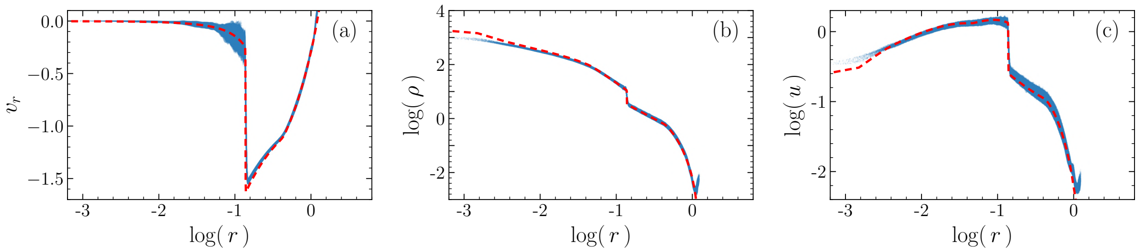

Figure 2.15: Evrard collapse with a resolution of $\sim 10^7$ SPH particles at time $t = 0.8$, simulated using REMIX. Plots show (a) radial velocity, v_r, (b) density, ρ, and (c) specific internal energy, u, plotted against radial distance from the cloud centre, r. Individual particles are plotted in blue, and the dashed red line shows a reference solution from a high-resolution grid code simulation (Borrow et al., 2022).

act to smooth this noise over time.

The Evrard collapse is captured well by REMIX, as shown in Fig. 2.15. We observe sharp shocks and evolution that closely follows the reference solution. The scatter in internal energy around the reference solution could be reduced by increasing the strength of artificial diffusion of internal energy, through choices of a_u and b_u. However, we choose to maintain a conservative approach to artificial diffusion so as not to deviate far from the thermodynamically consistent basis of our equations of motion. We therefore judge this amount of scatter to be sufficiently small. There is less scatter in density than in internal energy, since the normalising term is also contributing to smoothing the density.

At the vacuum boundary, we see a slight upturn in density and internal energy. Since divergence estimates in the evolution of these quantities revert to using kernels that are normalised to the continuum at vacuum boundaries, bulk expansion at vacuum boundaries may be underestimated. This is because for a region of locally isotropic expanding gas, a spherically symmetric kernel that is sampled by diverging particles in only approximately half its volume will underestimate the local velocity divergence. We note however that the logarithmic scales in Fig. 2.15 perhaps overemphasise the upturning features in terms of their importance in a typical science application.

The evolution of energy in Evrard collapse simulations is shown in Fig. 2.16. The exchange of energy between different forms is closely aligned between REMIX and tSPH simulations, as demonstrated in Fig. 2.16(a). These curves are consistent

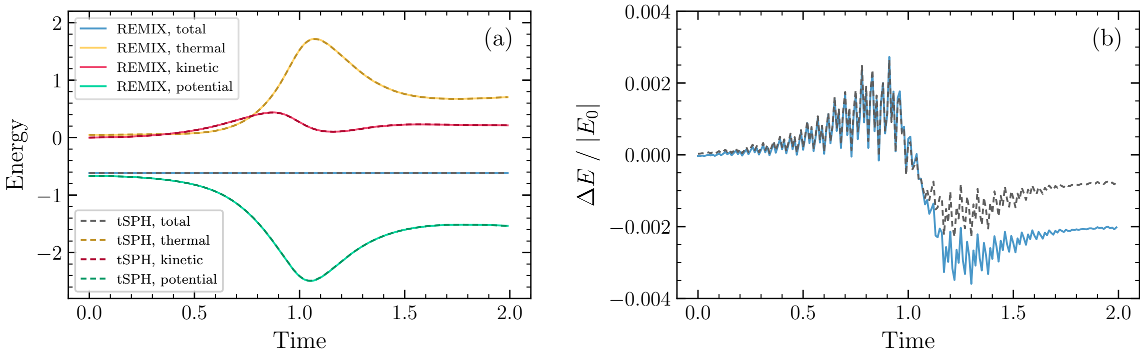

Figure 2.16: Energy evolution in the Evrard collapse. (a) Different forms of energy, and (b) fractional deviation of total energy from its initial value, are shown as functions of time for tSPH and REMIX simulations, both with a resolution of $\sim 10^7$ particles. The line colours and styles in (b) match the total energy in (a).

with those shown in Fig. 42 of Springel (2010b). The fractional deviation of the total energy from its initial value is plotted in Fig. 2.16(b). Both REMIX and tSPH are constructed to explicitly conserve energy. Fluctuations of total energy are of the same order of magnitude in both cases, with REMIX showing variations of less than 0.4% during the simulation. Small deviations of energy of this size are expected for SPH schems with non-reversible timesteps, even in formulations like ours, whose governing equations are explicitly conservative.

2.3.9 Planets in hydrostatic equilibrium

Planets in hydrostatic equilibrium offer a test scenario to probe the interaction of our hydrodynamic methods with gravity in the context of a layered, multi-EoS structure with a free surface. This also acts as a useful validation test for potential science applications. We also use this test to illustrate the importance of the inclusion of both the vacuum boundary treatment and the density evolution kernel normalising term in the REMIX scheme.

We apply our methods to an Earth-like and a Jupiter-like planet. The Earth-like planet represents a case in which materials have small variations of density within layers. The Jupiter-like case represents a scenario with relatively steep gradients of densities within material layers. This acts to assess stability against error-driven instabilities that can form due to these density gradients. These are only "Earth-like" and "Jupiter-like" because they are based on initial conditions for planetary

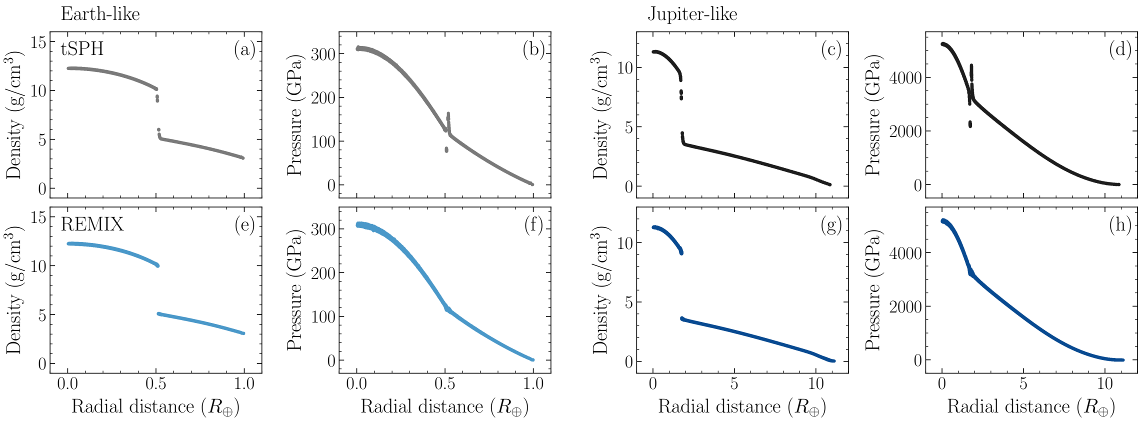

Figure 2.17: Radial profiles from simulations of an Earth-like (a, b, e, f) and a Jupiter-like (c, d, g, h) planet at time $t = 10{,}000\,\mathrm{s}$, simulated using tSPH (a–d) and REMIX (e–h). Particle densities and pressures are plotted against radial distance from the centre of the planet. REMIX corrects density discontinuities in simulations of planets in hydrostatic equilibrium.

giant impact simulations, which resemble the present day planets after the impact (Kegerreis et al., 2022). However, unlike in typical pre-impact "settling" simulations, where particle entropy can be fixed to prevent viscous heating (Kegerreis et al., 2020), here we use the full REMIX scheme with no modifications.

Initial hydrostatic equilibrium profiles and SPH particle placements are calculated using the publicly available code WoMa (Kegerreis et al., 2019; Ruiz-Bonilla et al., 2021). The Earth-like planet is constructed to satisfy the following conditions: two adiabatic layers consisting of a core of mass $0.27\,M_\oplus$, where $M_\oplus = 5.97\times10^{24}$ kg, represented by particles with the ANEOS $\mathrm{Fe_{85}Si_{15}}$ (iron) EoS and a mantle of mass $0.62\,M_\oplus$ with ANEOS forsterite (rock) (Stewart et al., 2020); a surface pressure and temperature of $P_\mathrm{s} = 1\times10^5\,\mathrm{Pa}$ and $T_\mathrm{s} = 2000\,\mathrm{K}$. The Jupiter-like planet is constructed to satisfy the following conditions: two adiabatic layers consisting of a core of mass $10\,M_\oplus$ with the AQUA (ice) EoS (Haldemann et al., 2020), and a hydrogen–helium (Chabrier & Debras, 2021) envelope of mass $298\,M_\oplus$; a surface pressure and temperature of $P_\mathrm{s} = 1\times10^5\,\mathrm{Pa}$ and $T_\mathrm{s} = 165\,\mathrm{K}$. In all our simulations, planets each consist of $\sim10^7$ equal mass particles.

Fig. 2.17 shows radial density and pressure profiles of the two planets at a time $t = 10{,}000\,\mathrm{s}$. We show profiles from simulations using both tSPH (a–d) and REMIX (e–h). The smoothing of density discontinuities in tSPH, and the corresponding

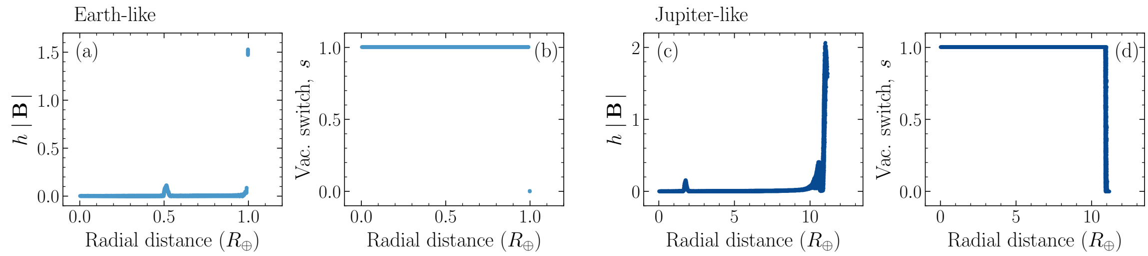

Figure 2.18: Identification of planetary vacuum boundaries in REMIX simulations. Plots correspond to Earth-like (a, b) and Jupiter-like (c, d) planets at time $t = 10{,}000\,\mathrm{s}$. We plot the quantity $h\,|\mathbf{B}|$ for individual particles, which is used in our vacuum boundary switch (Eqn. 2.21), and the vacuum boundary switch, s, itself.

pressure discontinuities, are clearly visible. We note that at this time particles have evolved to take up more relaxed positions, so density smoothing at the material interface, and in particularly at the vacuum boundary, are less extreme than in the initial condition configuration (Appendix H). However, these relaxed configurations typically yield large gaps between the different-density layers, so are a result of the surface tension rather than an indication that surface tension is reduced as the system relaxes. With REMIX, density discontinuities remain sharp, and pressures at the material boundaries are close to continuous.

We use this test to motivate the inclusion of the vacuum boundary treatment, detailed in §2.2.4. The vacuum boundary switch is able to accurately identify free surfaces based on $h_i\,|\mathbf{B_i}|$ (Eqn. 2.21), as shown in Fig. 2.18 for the two example planets. In the Earth-like planet the outermost particles remain in an undisturbed shell, which all get identified as the vacuum boundary and no interior particles are flagged. In the Jupiter-like planet, however, the envelope density drops far lower before the outer edge, leading to steep local changes in density near the vacuum boundary. This leads to error-driven particle motion that disturbs the initial particle shells, demonstrated in the identification of the vacuum boundary by the switch function, which in this case extends smoothly to particles near the surface that are no longer neatly ordered in shells.

In Fig. 2.19(a) and (d), we show a cross-section and density profile from a REMIX simulation of the Jupiter-like planet at $t = 20{,}000\,\mathrm{s}$ without the vacuum boundary treatment in the kernel construction, although still included in the normalising term. In this case, linear-order reproducing kernels are used without modification for all

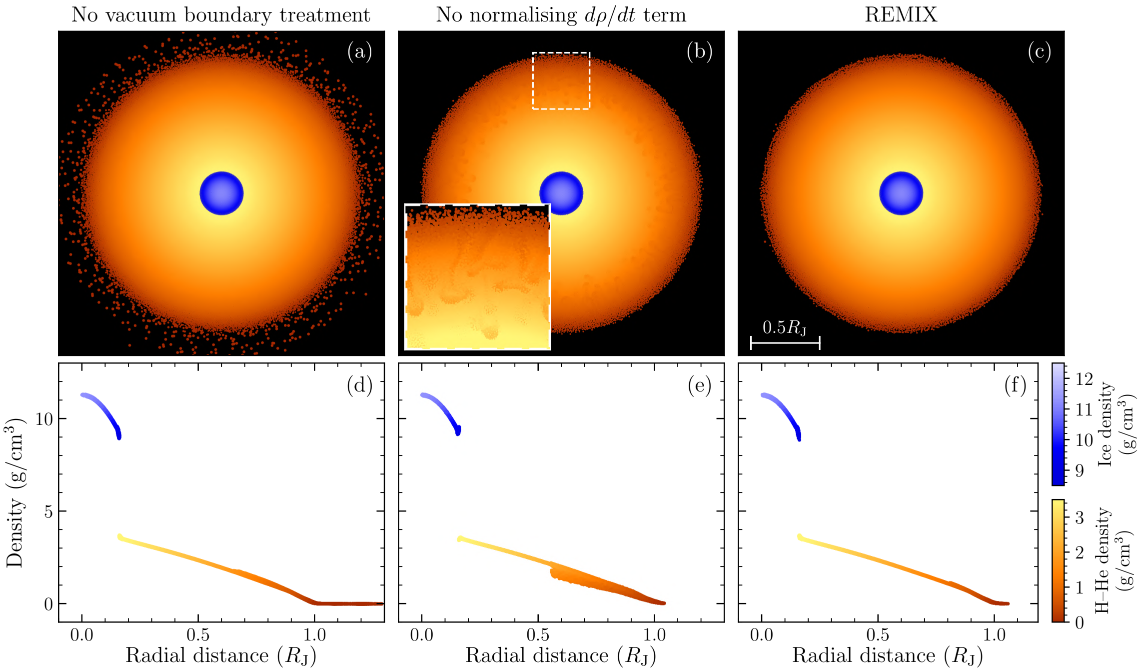

Figure 2.19: Effect of the vacuum boundary treatment and the normalising term in simulations of a Jupiter-like planet in hydrostatic equilibrium. We plot snapshots (a–c) and radial density profiles (d–f) at time $t = 20{,}000\,\mathrm{s}$. Columns show simulations with: the REMIX scheme but without the inclusion of the vacuum boundary treatment (a, d); the REMIX scheme but without the inclusion of the normalising term in the density evolution (b, e); and the full REMIX scheme with no modification (c, f). Particles are coloured by material and density. The inset in (b) shows a magnified view of instabilities forming near the vacuum boundary when the normalising term is not included. The colour scale of the inset has been slightly tweaked to increase the contrast around the instabilities. The evolution of these instabilities, without tweaked colours, is shown in Fig. H.2.

particles across the simulation, including particles near the free surface. Particles become unstable at the vacuum boundary, despite being set up to satisfy hydrostatic equilibrium, because bad estimates of pressure gradients lead particles to stream out from the surface. Similar behaviour is observed in equivalent simulations of the Earth-like planet.

In Fig. 2.19(b) and (e), we show similar results from a simulation with REMIX, but without the normalising term in the density evolution equation (§2.2.7). Here we see error-driven instabilities forming near the vacuum boundary. In Appendix H, we show the continued evolution of these instabilities to demonstrate how they continue to disturb the profile of the planet. These low density plumes that fall towards the planet's centre have a high local number density of particles, so should

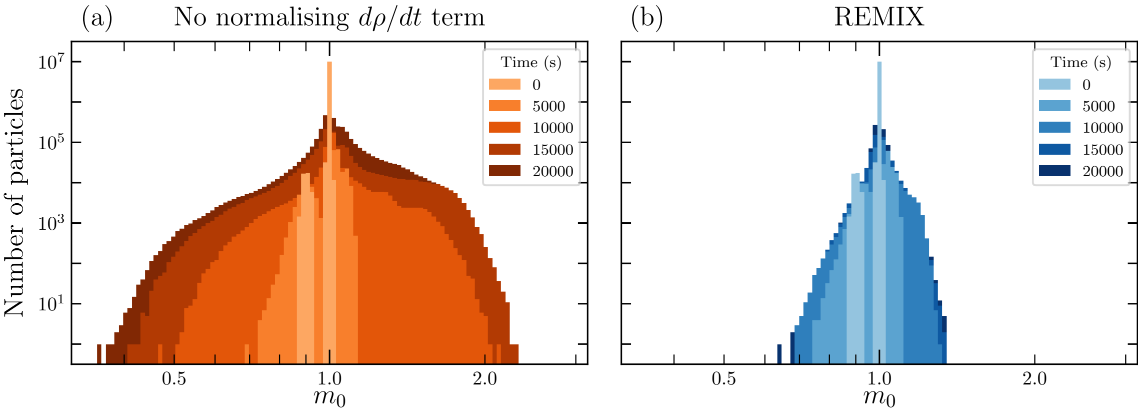

Figure 2.20: The evolution of the distribution of the m_0 kernel geometric moment in simulations of a Jupiter-like planet. Plots show results from simulations with a resolution of 10^7 particles using the REMIX scheme both (a) without and (b) with the normalising term in the density evolution. Only particles with a vacuum boundary switch $s_i > 0.9$ are plotted, to isolate particles that should have $m_0 \approx 1$ and filter out spikes near $m_0 = 0.5$.

have higher densities. This disconnect between the particle density and the local distribution of mass in the simulation volume leads to a positive feedback effect, in which the falling plumes continue to accumulate particles without the density of particles evolving to reflect this, further driving the discontinuity downwards. In the full REMIX scheme, for which similar plots are shown in Fig. 2.19(c) and (f), we re-associate the evolved density to the mass distribution in simulation volume by the inclusion of the kernel normalising term, which prevents the formation of these instabilities. We show the direct effect of the normalising term in these simulations in Fig. 2.20. As in the KHI examples in Fig. 2.7, here we see how the normalising term acts to tie particle volume elements to the local distribution of particle masses. The signal velocity of the normalising term means that correction occurs over the timescale of particle motion. Therefore, particles are able to readjust to react to changes in density due to this term, while it still acts as an effective correction to accumulation of error, accumulated over timescales set by the velocity divergence estimate in the equations of motion. In regions where densities represent the local distribution of particle masses well, the normalising term has little effect on the hydrodynamics.

2.4 Conclusions

In this Chapter, I have presented a new formulation of smoothed particle hydrodynamics (SPH), REMIX (‘Reduced Error MIXing’), that combines several novel and recently developed methods to address the well-known shortcomings of the traditional SPH formalism at density discontinuities. By directly targeting sources of kernel smoothing error and discretisation error, this scheme dramatically reduces numerical effects that can otherwise lead to spurious surface tension-like effects and inhibit mixing. We demonstrate its effectiveness using 3D hydrodynamic tests in a broad range of scenarios and regimes. In addition to standard tests, REMIX can handle boundaries between dissimilar, stiff materials, and the particularly challenging case of density discontinuities in simulations with equal mass particles — where both smoothing and discretisation errors are considerable.

The REMIX SPH scheme is based on thermodynamically consistent, conservative equations of motion, with free functions chosen to limit zeroth-order error. We use an evolved density estimate to avoid the kernel smoothing error in the standard SPH integral density estimate. To avoid potential accumulation of error in the evolved density estimate, such that densities would be no longer representative of the distribution of particle masses in the simulation volume, we introduce a new “kernel normalising term”. Additionally, artificial diffusion, which is weak outside shocks, helps to smooth out accumulated noise in both particle densities and internal energies. To reduce discretisation error, we use linear-order reproducing kernels in the equations of motion. Since kernel densities are evolved in time, particle volume elements are not *instantaneously* tied to simulation volume, despite the normalising term in the density evolution. Therefore, normalising the kernel to particle volume elements is an important step in calculating appropriate gradient estimates in the equations of motion. We introduce grad-h terms to the kernels, adding completeness to the construction. Additionally, we present a novel method that identifies free surfaces and reverts our kernels to standard spherically symmetric functions, normalised to the continuum to appropriately capture vacuum boundaries. We also use advanced artificial viscosity and diffusion schemes with linear reconstruction of quantities to particle midpoints, and a set of novel improvements to effectively switch between treatments for shock-capturing under compression and noise-smoothing in shearing regions.

REMIX shows a range of improvements compared with traditional SPH formulations, as we examined here with an extensive set of test cases. Our generalised error-reduction approach greatly improves the treatment of both static density discontinuities, as seen in the 3D square test with equal mass particles, as well as mixing and instability growth at evolving interfaces within a single ideal gas and between multiple materials, as demonstrated in fluid instability tests. This is done without need for a material-dependent approach in volume elements or density estimates, and without applying targeted corrections to material boundaries or need for bespoke choices of particle masses. REMIX is able to capture shocks with reduced particle noise, as seen in the Sod shock tube, and can effectively simulate a system with gravity and emerging shocks where large amounts of energy are exchanged between different forms, as demonstrated in the Evrard collapse. Many aspects of REMIX combine to allow us to improve simulations of planetary bodies, including the evolved density estimate that corrects smoothing of density discontinuities; the vacuum boundary treatment that extends our methods to be able to deal with free surfaces; and the density evolution normalising term that ensures that particle densities are tied to the local distribution of masses.

CHAPTER 3

Giant impacts onto Jupiter

3.1 Introduction

Measurements of Jupiter's gravitational moments by the Juno spacecraft have led to models of the planet's interior that suggest the existence of a dilute core: an extended compositional gradient between Jupiter's central core of heavy elements and its hydrogen–helium envelope (Debras & Chabrier, 2019; Howard et al., 2023; Miguel et al., 2022; Militzer et al., 2022; Nettelmann, 2017; Vazan et al., 2018; Wahl et al., 2017). This is inconsistent with traditional giant planet formation models that predict a differentiated internal structure (Müller et al., 2020). With kronoseismology suggesting that Saturn also has a dilute core (Mankovich & Fuller, 2021), understanding the processes that govern the formation of such compositional gradients would provide key insights into the evolution of giant planets and planetary systems.

Several mechanisms have been proposed to explain Jupiter's dilute core (Helled et al., 2022). An extended planetesimal-dominated accretion phase could lead to the dilute core being in place prior to runaway gas accretion (Venturini & Helled, 2020). Alternatively, convective processes could gradually erode a differentiated core until it reaches a mixed state (Moll et al., 2017). Liu et al. (2019) (hereafter L19) proposed a giant impact as an alternative mechanism. Giant impacts have previously been suggested as a mechanism that can significantly affect the internal structure of giant planets in various ways (Li et al., 2010; Liu et al., 2015).

The head-on impact simulation of L19 presented the disruption of a differentiated core into a well-mixed, diluted state with a heavy-element fraction in the centre of the planet of $Z < 0.5$, by a 10 $M_\oplus$ impactor. Their simulations with a larger impact angle or a smaller impactor mass did not produce a dilute core. The hydrodynamic simulations of L19 were carried out using the adaptive mesh code FLASH (Fryxell et al., 2000). By modelling the subsequent thermodynamic evolution, L19 found that this compositional gradient could persist for Gyr timescales until the present day. However, overmixing in regions of large bulk motion relative to the stationary grid points is a typical shortcoming of Eulerian methods. This spurious diffusion arises from the advection terms necessary in this non-Lagrangian method (Robertson et al., 2010; Springel, 2010b). Additionally, the accuracy of the treatment of self-gravity is sensitive to choices made in the multipole approximation of the gravitational potential (Couch et al., 2013), and L19 used idealised equations of state that do not capture the complexities of metallic hydrogen within Jupiter's deep interior (Chabrier et al., 2019). Therefore, it would be valuable to investigate the equivalent impact scenario using fundamentally different modelling approaches, to assess the potential sensitivity of dilute core production to the specifics of the numerical methods employed.

Lagrangian hydrodynamic methods, where Galilean invariance is maintained, do not experience this kind of artificial mixing from advection through grid points, since interpolation points move with the fluid velocity. In particular, smoothed particle hydrodynamics (SPH) is widely used for simulations of giant impacts since it: inherently tracks the evolution of fluid element trajectories and thermodynamics; is able to deal with vacuum regions and evolving free surfaces efficiently; offers geometry-independent adaptive resolution; and couples elegantly with gravity solvers (Gingold & Monaghan, 1977; Lucy, 1977). In traditional SPH formulations, however, mixing at density discontinuities is typically suppressed by spurious surface tension-like effects (Agertz et al., 2007). These artificial effects are considerable – and challenging to remedy – at boundaries between dissimilar, stiff materials for which more significantly erroneous estimates of fluid pressure suppress mixing more strongly (Ruiz-Bonilla et al., 2021). Therefore, a more advanced SPH construction that directly addresses these known sources of error is needed to utilise the benefits of the SPH formulation to reliably investigate dilute core formation in giant impact simulations, where material mixing is the key physical mechanism of interest.

REMIX is an advanced SPH scheme designed to directly address the sources of numerical error that suppress mixing in SPH simulations (§2.2). The REMIX scheme incorporates a range of novel and recently developed improvements to traditional SPH formulations, and its construction is generalised to address sources of error independent of material type or equations of state. It has demonstrated significant improvements in treating both mixing and instability growth in simulations with materials and conditions representative of those in giant impact simulations (§2.2). REMIX is integrated into the open-source, state-of-the-art SWIFT code (Schaller et al., 2024), whose computational efficiency enables simulations of giant impacts to be performed at high resolutions (Kegerreis et al., 2022).

Here, we use REMIX SPH to investigate whether Jupiter's dilute core was formed by a giant impact, as well as traditional SPH comparisons. First, in §3.2 I describe the methods used to perform simulations and to construct initial conditions. In §3.3 we test REMIX in fluid instability simulations under conditions representative of Jupiter's deep interior. Then in §3.4 we use REMIX for giant impacts onto Jupiter. We carry out simulations at a range of impact speeds and angles (§3.4.2); different numerical resolutions (§3.4.3); and with pre-impact planet structures and equations of state set up to closely follow those of L19 (§3.4.4). I discuss our results in §3.5 and summarise our findings in §3.6.

3.2 Methods

3.2.1 REMIX smoothed particle hydrodynamics

As presented fully in §2.2, REMIX is an SPH formulation designed to address key sources of error that suppress mixing and instability growth in traditional SPH simulations, particularly at density discontinuities. By adopting a generalised, material-independent approach, REMIX is able to not only improve the treatment of contact discontinuities within a single material but also handles well the more challenging case of interfaces between dissimilar, stiff materials. Like traditional SPH, REMIX inherently conserves mass, energy and momentum, is constructed from a basis of thermodynamic consistency, and is fully Lagrangian, ensuring Galilean invariance. REMIX has been extensively tested with standard and giant impact-relevant test scenarios, with full details in Chapter 2. Here, we summarise the primary features of REMIX and set up some additional test scenarios tailored directly to a Jupiter core-mixing context.

In traditional SPH ("tSPH") formulations used for applications in astrophysics, the fluid density at the positions of particles is estimated by kernel interpolation using an extended, Gaussian-like kernel function (Price, 2012). The standard SPH density estimate will smooth of the density field on kernel length scales (Violeau & Fonty, 2019). In regions where the density varies smoothly, this will be a minor effect. However, in sharply varying regions, and in particular at discontinuities in the underlying field, the reconstructed density field will inevitably be smoothed.

The effect of this kernel smoothing can be clearly seen at interfaces between different material layers in the pre-impact planets used for our giant impact simulations. Pre-impact planetary equilibrium profiles and the corresponding SPH particle placements are calculated using the publicly available WoMa and SEAGen codes (Kegerreis et al., 2019; Ruiz-Bonilla et al., 2021). Prior to impact simulations, additional adiabatic "settling" simulations are performed to allow particles to rearrange themselves towards an equilibrium configuration. In these simulations, particle entropies are fixed to their initial value to enforce adiabatic evolution. Settling simulations are carried out separately for each planet and are run for a simulation time

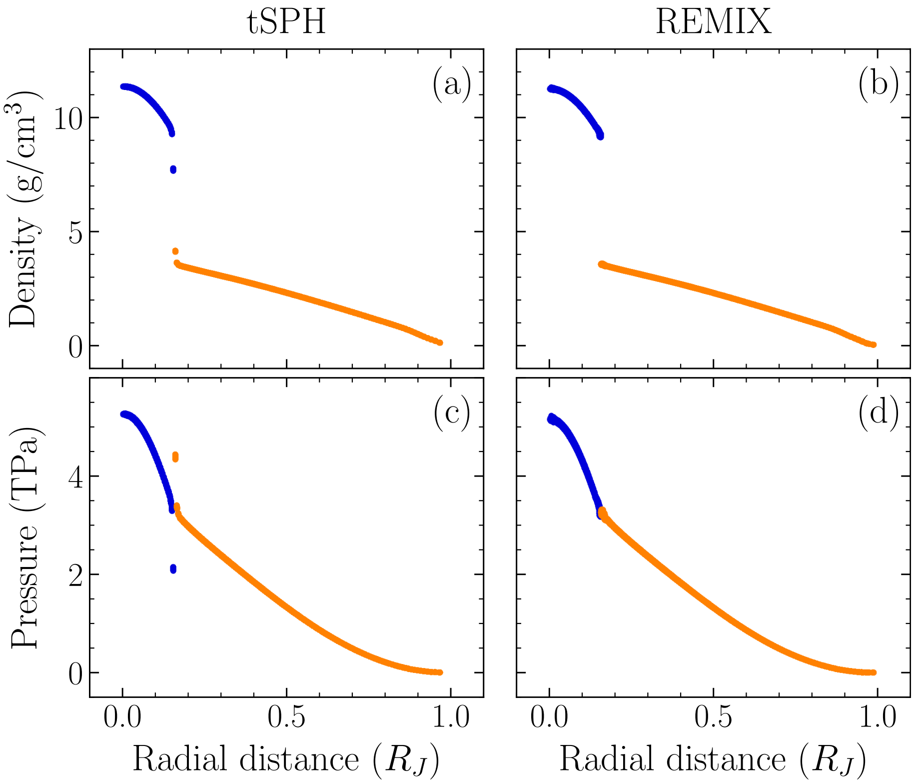

Figure 3.1: Radial profiles of density (a, b) and pressure (c, d) for a two-layer proto-Jupiter settling simulation at time $t = 5000$ s. Columns show profiles from simulations using a traditional SPH formulation ("tSPH") (a, c) and REMIX SPH (b, d). Individual particles are coloured by material type: blue for ice and orange for hydrogen–helium.

of 5000 s. REMIX reduces the errors that traditionally make calculations of particle accelerations sensitive to the local particle configuration. Therefore, the amount of particle motion in settling simulations is reduced.

The radial density and pressure profiles of a proto-Jupiter planet, to be used in our impact simulations, is shown in Fig. 3.1 from settling simulations using tSPH and REMIX. With tSPH the density field is smoothed by kernel interpolation, leading to diverging pressures at the core–envelope boundary that act as an artificial barrier to mixing across the interface. With REMIX, the density discontinuity stays sharp and the pressure remains continuous across the material boundary. To address kernel smoothing error, REMIX uses a differential form of the density estimate by which particle densities are evolved with time rather than recalculated from the instantaneous distribution of particle masses. All density discontinuities, including those at material boundaries and free surfaces, are not erroneously smoothed.

The smoothing error introduced by using an extended kernel function combines with the error introduced by the discretisation of the underlying fluid into a finite set of particles (Price, 2012; Spreng et al., 2020). To deal with discretisation error and mitigate the accumulation of error over time in evolved densities and internal energies, REMIX uses: linear-order reproducing kernels (Frontiere et al., 2017) that treat free surfaces as vacuum boundaries; a choice of free functions in the SPH equations of motion that limits discretisation error (Read et al., 2010); a kernel normalising term in the density evolution calculations; advanced formulations of artificial viscosity as well as artificial diffusion of internal energy and density between materials of the same type (§2.2).

All simulations presented here use particles of equal mass across the simulation, a case that was specifically considered in the validation of REMIX, and we employ the Wendland C^2 kernel with $\eta = 1.487$ to construct linear-order reproducing kernels (Dehnen & Aly, 2012; Wendland, 1995). Particles have approximately 100 neighbours within their kernel. REMIX has been developed with computational efficiency in mind and therefore, as demonstrated in the simulations presented here, can be used in simulations at cutting-edge resolutions for giant impact simulations.

3.2.2 Equations of state

The equations of state (EoS) characterise the thermodynamic behaviour of a material. In the SPH simulations presented here, the EoS are used to calculate pressures and sound speeds, from densities and internal energies. These quantities are then used both to directly evolve the simulated fluid and in calculations of timestep durations.

Models of Jupiter's internal structure are sensitive to uncertainties in the hydrogen–helium EoS used to calculate the planet's envelope profiles (Howard et al., 2023; Mazevet et al., 2022; Miguel et al., 2016). As such, much work is ongoing to create equations of state that accurately reproduce the behaviour of hydrogen–helium at the extreme densities and pressures within Jupiter (Becker et al., 2014; Chabrier & Debras, 2021; Chabrier et al., 2019; Militzer & Hubbard, 2013; Saumon et al., 1995). For simulations of giant impacts onto Jupiter and hydrodynamic tests using Jupiter-like materials, we use the Chabrier & Debras (2021) hydrogen–helium EoS (hereafter CD21 H–He), with a helium mass fraction of $Y = 0.245$ (Chabrier et al., 2019). For simulations of impacts onto Jupiter aiming to reproduce directly the initial conditions of L19, we use an ideal gas with adiabatic index $\gamma = 2$.

For heavy elements we use the AQUA EoS (Haldemann et al., 2020) to represent ice and the ANEOS forsterite EoS (Stewart et al., 2020) for rocky material. For direct L19 comparison simulations, we use Tillotson ice and granite (Melosh, 1989).

3.2.3 Impact initial conditions

For the majority of our simulations, we use a two-layer proto-Jupiter with an ice core and H–He envelope, and a single-layer ice impactor. For simulations set up to most closely match the initial conditions of the simulations of L19, we use 3-layer pre-impact planets with layers of rock, ice, and gas for both target and impactor. The choice to focus on impacts between planets with a reduced number of layers is made to further reduce any potential barriers to mixing. In all of our simulations, we follow L19's scenario and use an impactor with a total mass of 10 $M_\oplus$ and a proto-Jupiter of total mass 308 $M_\oplus$, with core mass of ∼10 $M_\oplus$, where $M_\oplus = 5.972 \times 10^{24}$ kg. The total mass of the system is therefore the present-day mass of Jupiter, $M_J = 1.898 \times 10^{27}$ kg. Simulations with three-layer planets have a slightly more massive core of 11.6 $M_\oplus$ to give equilibrium profiles that more closely match those of L19, although we find that changes to the initial profiles do not significantly affect the evolution of the impact. H–He layers are chosen to be adiabatic with surface temperatures (defined by where $P = 1$ bar) of 165 K for proto-Jupiters and 500 K for impactors. For two-layer proto-Jupiters and single-layer impactors, the ice layer is also chosen to be adiabatic with an impactor surface temperature of 200 K. For three-layer planets, the temperature–density relation of heavy-element layers is chosen somewhat arbitrarily, to attempt to match the radii of material interfaces of the simulations of L19.

All impact simulations are performed in 3D. They are set up 1 h prior to impact, defined as the planets' individual centres of mass reaching the summed distance of their initial radii, such that the shapes of the planets are allowed to realistically distort under tidal forces. Most impacts performed here are head-on, following the simulation of L19 that produced a dilute core. At Jupiter's orbital distance from the Sun, we expect the peculiar velocity of the impactor to be small compared with the mutual escape speed, $v_{\rm esc} = 54$ km s^{-1}, and therefore we simulate most impacts with an impact velocity of $v = v_{\rm esc}$. Note that L19 simulate impacts with $v = 46$ km s^{-1}, at the point of impact. Our impact parameter space exploration includes speeds as low as $v = 40.5$ km s^{-1} to test the potential implications of this choice.

We carry out a suite of simulations to systematically probe the effect of impact speed ($v = 0.75$, 1.0, 1.5 $v_{\rm esc}$); impact angle (with impact parameter $b = 0.0$, 0.2, 0.4, 0.6); and numerical resolution (with particle number $N = 10^5$–$10^{8.5}$ in logarithmic steps of $10^{0.5}$). These simulations are based on our fiducial simulation that uses planets with a reduced number of layers, is head-on, is at the mutual escape velocity, and has resolution 10^7. We also carry out simulations to replicate the impact of L19 even more closely, with three-layer bodies using the EoS used in their simulations, as well as with the more sophisticated EoS detailed above.

3.2.4 Measures of material mixing

Parameterising material mixing will enable us to quantitatively describe the degree to which core material may be diluted throughtout the impact. We measure the state of mixing in our impact simulations using two parameters: the local heavy-element mass fraction, Z, and the total mass of mixed material across the simulation, $M_{\rm mix}$. These quantities describe the local and global state of material mixing respectively. In our simulations, mixing is treated at the particle scale and not below. Therefore, the material of each particle remains fixed for the duration of the simulation. To estimate mixing, we therefore calculate these quantities as weighted estimates based on the localised distributions of particle material types.

We use kernel interpolation to estimate the local heavy-element mass fraction. The quantity $\bar{Z}$ is calculated based on weighted contributions from nearby SPH particles. This parameter describes the fraction of local mass that is represented by heavy-element SPH particles, such that $\bar{Z} = 0$ in regions where no local particles are heavy elements and $\bar{Z} = 1$ where they all are. We estimate $\bar{Z}$ at the positions of SPH particles by

$$\bar{Z}_i \equiv \frac{\sum_j \zeta_j \, m_j \, W_{ij} V_j}{\sum_j m_j \, W_{ij} V_j} \; . \tag{3.1}$$

Here subscripts denote quantities either sampled at the position of, or associated with, a particle i or its neighbouring particles j. Sums are approximations of integrals over discrete volume elements $V_j = m_j/\rho_j$, where m_j and ρ_j are particle masses and densities. The kernel function $W_{ij} \equiv W(\mathbf{r}_{ij},\, h_i)$ contributes weighting based on the particle separation $\mathbf{r}_{ij}$ and is characterised by the smoothing length h_i. The parameter ζ takes the value 1 if particle j's material represents heavy elements, and is 0 otherwise. We use the spherically symmetric Wendland C^2 kernel function for these calculations (Wendland, 1995). This is essentially a traditional-SPH measure of local Z for each particle. We use this rather than the linear-order reproducing kernels used in REMIX, since they would provide a less simple and less method-independent measure of the mixing.

To estimate the total mass of mixed material in our simulations we first define what constitutes a mixed state. Since each particle retains its material for the duration of the simulation, we determine that a particle with neighbours of different material-types only – with no neighbours of its own type – is in a maximally mixed state. For a particle i we estimate the local mass fraction of particle i's own material, similarly to $\bar{Z}$, by

$$\bar{w}_i \equiv \frac{\sum_j \kappa_{ij} \, m_j \, W_{ij} V_j}{\sum_j m_j \, W_{ij} V_j} \; , \tag{3.2}$$

where $\kappa_{ij} = 1$ for particle pairs of the same material and $\kappa_{ij} = 0$ otherwise. We note that, unlike $\bar{Z}$, the value of $\bar{w}$ will never reach 0 because of the contribution of i itself in this calculation. We estimate that the contribution to m_i from materials *different* from that of i to be $m_{\mathrm{mix},\,i} \equiv (1 - \bar{w}_i)\; m_i$. The total mixed mass in the simulation is then given by

$$M_{\mathrm{mix}} \equiv \sum_i m_{\mathrm{mix},\,i} = \sum_i \, (1 - \bar{w}_i)\; m_i \; , \tag{3.3}$$

where we sum over all simulation particles.

We note that both $\bar{Z}$ and $M_{\rm mix}$ will be spatially smoothed on the scale of the smoothing length, since they are calculated by interpolation using an extended kernel. Therefore material near sharp material interfaces will be measured as mixed even if particles of different materials have not crossed the interface.

3.3 Fluid instabilities and mixing

Before running the primary impact simulations, we first test REMIX in simulations of Kelvin–Helmholtz and Rayleigh–Taylor instabilities with materials and conditions representative of material interfaces in giant impacts onto Jupiter. Although no converged reference solutions exist for these scenarios, demonstrating that REMIX alleviates the purely numerical effects at the material interface will verify that the material-independent improvements of REMIX apply in this regime, where core-material and metallic hydrogen have been predicted to be miscible (Wilson & Militzer, 2011, 2012). We carry out fluid instability simulations in 3D and with particles of equal mass across the simulation to validate our hydrodynamic treatment for our impact simulations, as done in §2.2 for similar but not Jupiter-specific tests.

3.3.1 Kelvin–Helmholtz instability

The Kelvin–Helmholtz instability (KHI) arises as perturbations at shearing fluid interfaces grow to form spiralling vortices (Chandrasekhar, 1961). We examine the growth of the KHI between layers of ice and H–He at conditions representative of Jupiter's deep interior. In our simulations these materials are treated as inviscid fluids and so, since the growth of the instability is predominantly inertial, we expect a qualitatively similar evolution to analogous, well-studied ideal gas simulations (Frontiere et al., 2017; McNally et al., 2012; Price, 2008; Robertson et al., 2010; Rosswog, 2020). We characterise the growth of a mode of wavelength λ by the timescale

$$\tau_{\rm KH} = \frac{(\rho_1 + \rho_2)\ \lambda}{\sqrt{\rho_1 \rho_2}\,|v_1 - v_2|}\ , \tag{3.4}$$

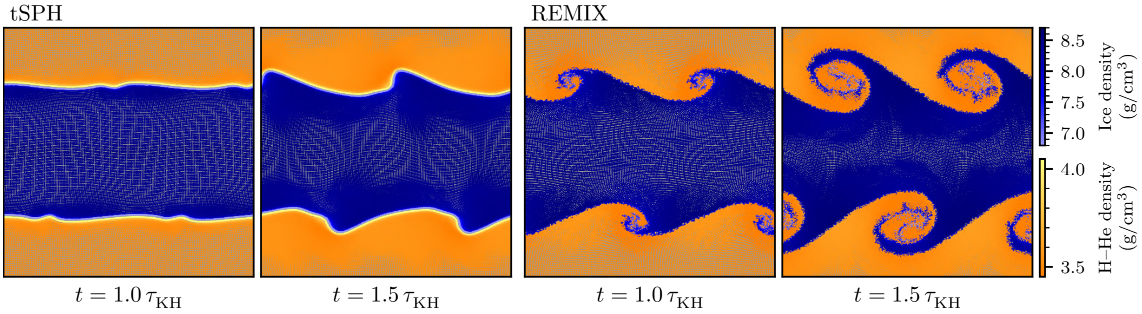

Figure 3.2: Kelvin–Helmholtz instability growth with materials and conditions representative of a pre-impact Jupiter's core–envelope interface. Snapshots show two times from simulations using a traditional SPH formulation ("tSPH") and REMIX. Individual particles are plotted on a grey background and coloured by their material type and density. Particles at all z are plotted, so the grey background is visible in regions that have maintained their grid alignment in z from the initial conditions.

where ρ_1 and ρ_2 are the densities in regions separated by the shearing interface and $|v_1 - v_2|$ is their relative speed (Price, 2008).

Initial conditions are constructed similarly to those of §2.2. H–He particles are initialised in a 3D cubic lattice in a periodic domain with $128 \times 128 \times 18$ particles in the x, y, z directions. The size of the simulation domain in x and y is 1 $R_{\rm J}$, where the radius of Jupiter is $R_{\rm J} = 69.9 \times 10^3$ km, and particle masses are chosen to give a density of $\rho_1 = 3.5$ g cm^{-3}. A region occupying the central half of the domain in y and spanning the full domain in x and z is replaced by a region of higher-density ice at, $\rho_2 = 8.43$ g cm^{-3}. These densities are chosen to correspond to the densities at the core–envelope interface in the pre-impact proto-Jupiter, as plotted in Fig. 3.1. Since we use particles of equal mass across the simulation, the cubic lattice of ice particles is initialised with a smaller grid-spacing. The particle configurations in both regions are constructed to maintain their grid-spacing across boundaries of the periodic domain and for the two regions to be separated by the mean of the two grid-spacings at both interfaces. The two regions are initialised with relative speeds of $v_1 = -10^{-4}$ $R_{\rm J}$ s^{-1} and $v_2 = 10^{-4}$ $R_{\rm J}$ s^{-1}. A mode of wavelength $\lambda = 0.5$ $R_{\rm J}$ and of form $v_y = 0.01|v_1 - v_2| \sin(2\pi x/\lambda)$ seeds the instability. Initial internal energies are set such that the regions are in pressure equilibrium with $P(\rho, u) = 3.2 \times 10^{12}$ Pa. We note that the spurious smoothing of the density discontinuity in tSPH means that, unlike with REMIX, simulations with tSPH are not truly initialised in pressure

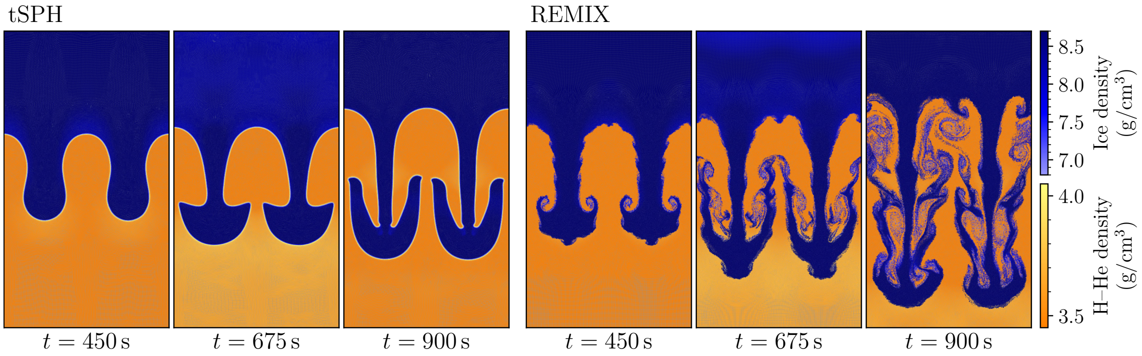

Figure 3.3: Rayleigh–Taylor instabilities for Jupiter-like materials and conditions. Snapshots show three times for simulations using a traditional SPH formulation ("tSPH") and REMIX. The regions of fixed boundary particles at the top and bottom of the simulations have been cropped from the figure; their positions and densities do not change.

equilibrium.

The evolution of the KHI with these initial conditions, from simulations using tSPH and REMIX, is shown in Fig. 3.2. The growth of the instability is clearly and strongly suppressed with tSPH: the characteristic spirals of the KHI do not form and particles are prevented from crossing the density discontinuity by spurious surface tension-like effects. REMIX directly addresses the sources of error that lead to these effects and so allows the instability to grow, and particles of different materials are able to intermix. The instability grows over a similar timescale, scaled by $\tau_{\rm KH}$, to the analogous KHI simulations with an ideal gas, and also those between Earth-like materials presented in §2.2.

3.3.2 Rayleigh–Taylor instability

The Rayleigh–Taylor instability (RTI) occurs due to the displacement of a high density fluid by a low density fluid (Chandrasekhar, 1961). We consider a gravity-driven case in which a region of dense ice sits above a region of H–He, initially in approximate hydrostatic equilibrium other than a small velocity perturbation. As in the KHI, spurious surface tension-like effects at the density discontinuity strongly suppress the growth of this instability in tSPH simulations.

Initial conditions are constructed similarly to those of §2.2. Particles are placed in a periodic simulation domain in two cubic lattices. The domain has dimensions of

0.5 R_J, 1 R_J in the x and y directions, with a thin 3.5×10^{-3} R_J domain size in the z dimension. The low density H–He region has $256 \times 256 \times 18$ particles with density $\rho_1 = 3.5$ g cm^{-3} and occupies the bottom half of the domain. The upper ice region is constructed to satisfy similar grid-spacing constraints as in the KHI simulation, with density $\rho_2 = 8.43$ g cm^{-3}. Particles in the top and bottom 0.05 R_J of the domain are fixed in place throughout the course of the simulation. Initial internal energies are set to satisfy hydrostatic equilibrium for a constant gravitational acceleration $g = -31.4$ m s^{-2}, and an interface pressure of $P_0 = 3.2 \times 10^{12}$ Pa, representative of the gravitational acceleration and pressure at the core–envelope boundary in the proto-Jupiter used for our fiducial giant impact simulations. Particles are initially at rest, other than an initial velocity perturbation that seeds the instability,

$$v_y(x, y) = \delta_y \left[1 + \cos\left(8\pi\left(x + 0.25\right)\right)\right]\left[1 + \cos\left(5\pi\left(y - 0.5\right)\right)\right] \tag{3.5}$$

in the region 0.3 $R_\mathrm{J} < y < 0.7$ R_J and $v_y = 0$ otherwise. We use a perturbation amplitude of $\delta_y = 0.025$ R_J s^{-1}.

The evolution of the RTI with these initial conditions is shown in Fig. 3.3, for simulations using tSPH and REMIX. In the tSPH simulation the RTI plumes grow slowly and material is prevented from crossing the interfaces. This is in contrast with the REMIX RTI, where we observe the unimpeded growth of both the primary and secondary instabilities. This leads to mixing across a range of length scales as particles are not artificially prevented from crossing the interface and instabilities grow to drive turbulent mixing.

The results of these KHI and RTI simulations demonstrate that REMIX is able to capture mixing and fluid instability growth in conditions representative of the giant impact simulations presented in the following section.

3.4 Giant impacts

Here we investigate dilute-core formation in simulations of: (1) head-on impacts carried out using both REMIX and traditional SPH (§3.4.1); (2) impacts at a range of speeds and angles (§3.4.2); (3) head-on impacts with different numerical resolu-

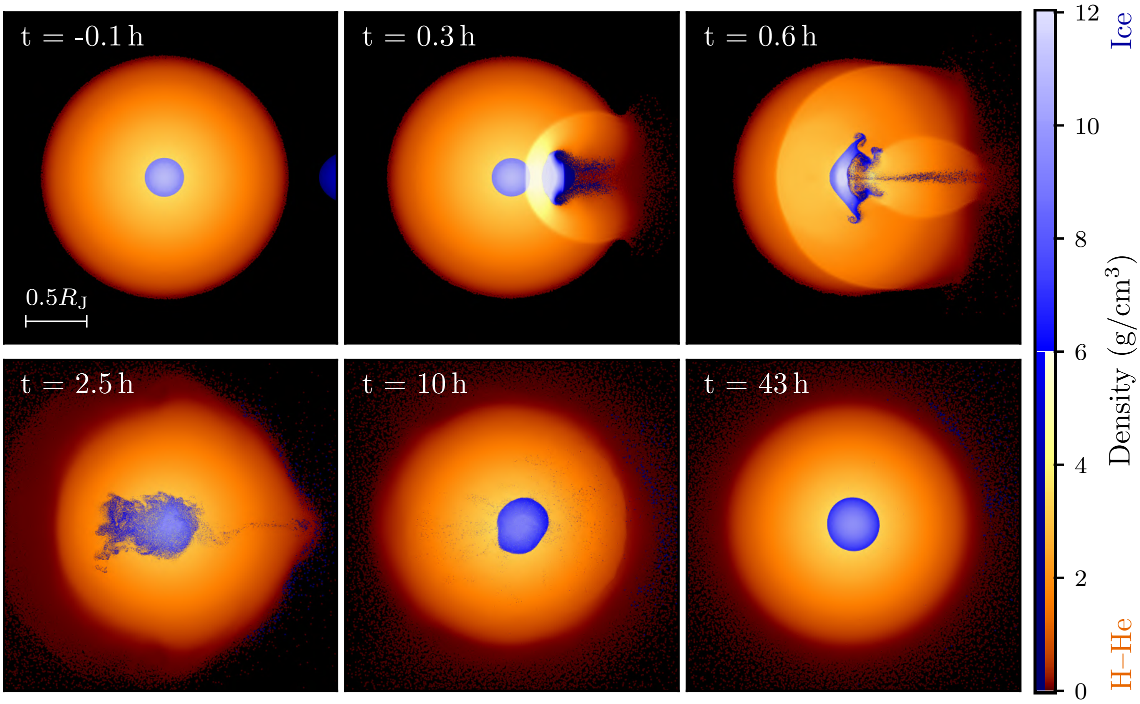

Figure 3.4: Snapshots from the fiducial, head-on impact onto Jupiter carried out using REMIX SPH. Individual SPH particles are plotted in cutaways from 3D simulations and are coloured by material-type and density.

tions (§3.4.3); (4) simulations with initial conditions set up to closely replicate those of L19, including an alternative version with more sophisticated equations of state (§3.4.4).

3.4.1 Fiducial scenario

As a basis for investigations of impact configuration and numerical resolution in later sections, we consider a fiducial scenario of the head-on impact between a 10 $M_\oplus$ impactor and a 308 $M_\oplus$ proto-Jupiter with a 10 $M_\oplus$ core, at the mutual escape speed of the two bodies. We choose to focus primarily on impacts with a two-layer proto-Jupiter with only layers of ice and H–He and a single-layer, ice impactor. We do this to deliberately reduce both the number of density discontinuities and the size of the core–envelope density contrasts in the initial conditions, to remove barriers to forming a dilute core in our simulations. We use the more realistic CD21 H–He and AQUA EoS for these simulations.

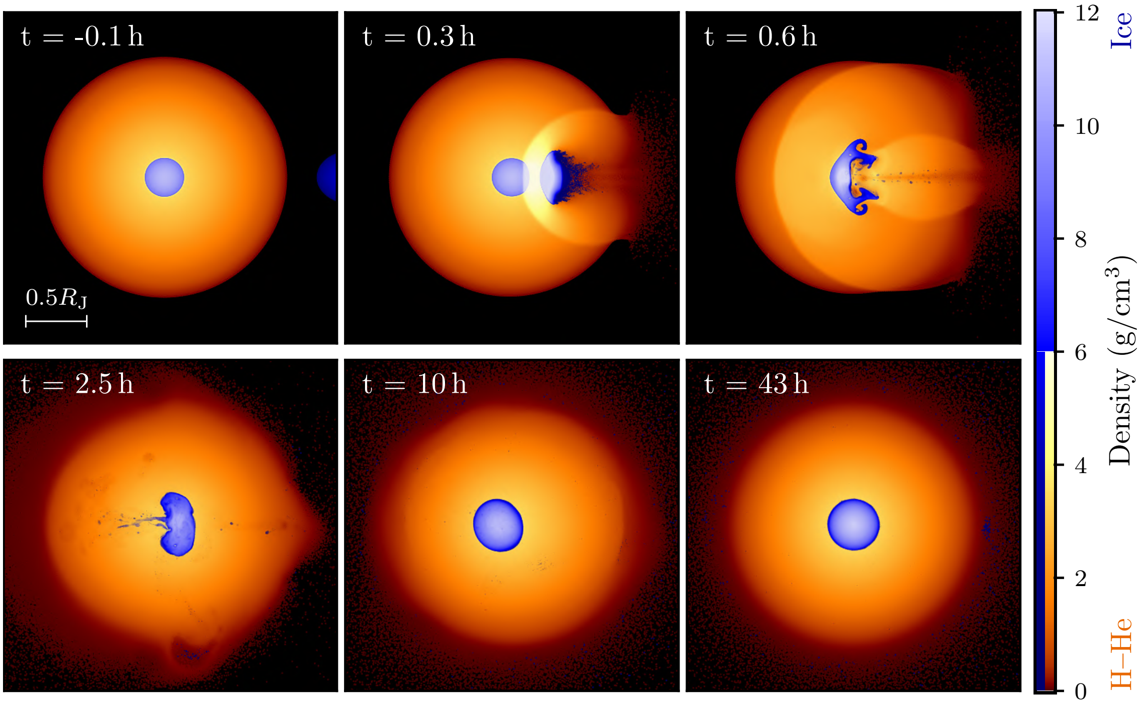

Figure 3.5: Snapshots from the fiducial, head-on impact onto Jupiter carried out using a traditional SPH formulation ("tSPH"). Individual SPH particles are plotted in cutaways from 3D simulations and are coloured by material-type and density.

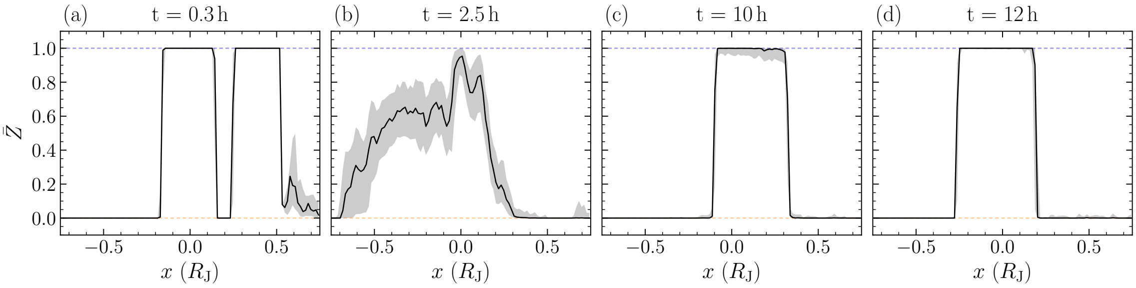

Figure 3.6: Profiles of localised heavy-element mass fraction, $\bar{Z}$, sampled in a thin, 0.05 R_J radius cylinder along the axis of head-on impact, x, from the fiducial REMIX simulation. Panels correspond to times: (a) immediately prior to core-disruption by the impactor; (b) when heavy elements are mixed with envelope material; (c) when the core has settled to form a discrete boundary, although still mixed with some envelope material; (d) when core- and envelope-material have largely separated. Black, solid lines show the median particle value in 100 bins along the plotted region. Grey shading spans percentiles such that the enclosed regions correspond to 68% of particles in each bin. The blue and orange dashed lines show heavy-element fractions of pure ice and hydrogen–helium respectively. The x axis is centred at the centre of mass of the system; deviations of the core position from $x = 0$ are due to post-impact oscillations.

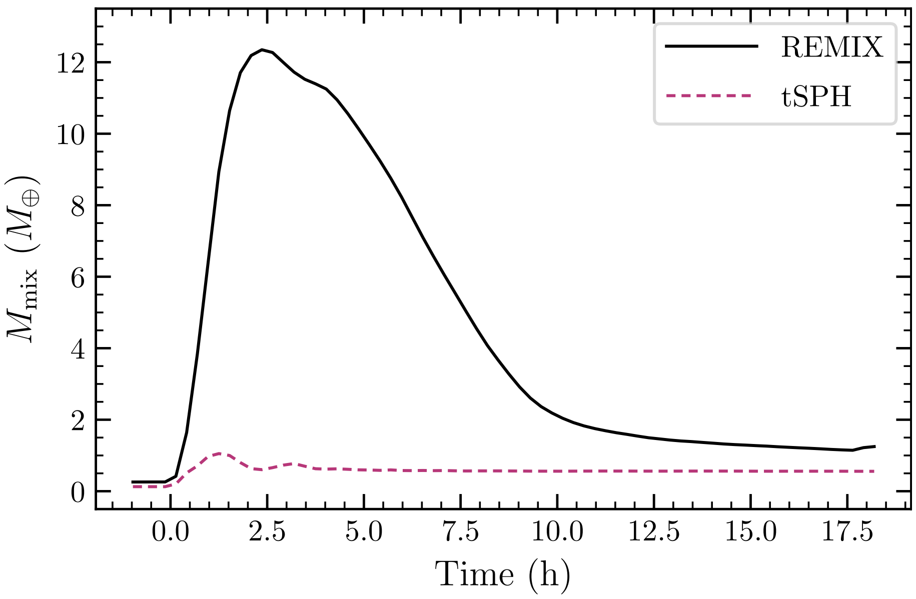

Figure 3.7: Evolution of the total mass of mixed material, $M_{\rm mix}$, in simulations of head-on impacts onto Jupiter. The black, solid line shows results from a REMIX simulation and the pink, dashed line from a tSPH simulation. The total mass of heavy elements is 20 $M_\oplus$ in these simulations.

We simulate the fiducial scenario using both tSPH and REMIX. In the REMIX simulation, ice particles can mix freely into the envelope, as seen in Fig. 3.4. The core reaches a temporarily mixed state, however, heavy elements rapidly settle under gravity to re-form a differentiated core over short timescales of ~10 h. The snapshots at 43 h show a later time where post-impact oscillations have dissipated. No dilute core is produced, even with the more realistic treatment of mixing in the REMIX scheme. In the tSPH simulation, spurious surface tension-like effects are strong, suppressing mixing of ice and H–He particles. Heavy elements remain in a largely cohesive mass throughout the simulation, which settles to form a core with a discrete interface between the two different materials, as shown in Fig. 3.5.

The evolution of the local heavy-element fraction (Eqn. 3.1) of the REMIX impact is shown in Fig. 3.6, for a thin cylinder aligned along the direction of the impact. At early times, $\bar{Z}$ reaches intermediate values as material mixes due to the erosion of the impactor as it travels through the envelope, as seen in Fig. 3.6(a), and due to the disruption of the core by the impact, the aftermath of which is shown in Fig. 3.6(b). At 10 h the core, not positioned at the centre of mass of the planet due to the oscillations, consists largely of heavy elements, and the core–envelope interface is already sharp. By 12 h the core is close to pure heavy elements.

The evolution of the total mixed mass (Eqn. 3.3) in these simulations is shown in Fig. 3.7. There is considerably more mixing with REMIX than with tSPH. In the REMIX simulation, mixing peaks at a time 2.5 h after impact, at which time the core- and impactor-material have been maximally disrupted and mixed with the H–He envelope. After this time, the mass of mixed material falls as the system settles under gravity and materials separate, with material being largely separated by ∼10 h. We carry out simulations until later times to allow the large, dynamical oscillations to dissipate, although for the majority of this time the boundary of the core is discrete and oscillations only affect its shape.

3.4.2 Impact speed and angle

Although no dilute core is produced in the head-on scenario of the previous section, this result could perhaps be sensitive to the speed and angle of the impact. The core of the planet is more likely to be disrupted in head-on or low-angle impacts, but impactor material may mix into the envelope more effectively by erosion in higher-angle impacts. Additionally, one might speculate that higher impact speeds may act to increase the initial material mixing, or lower speeds could lead to post-impact heavy-element distributions and internal energy profiles that are more stable to convection that may otherwise facilitate demixing of materials. Therefore, investigating a wide range of impact speeds and angles will enable us to examine the sensitivity of dilute core production to impact configuration.

For this parameter study, we use the same initial planetary bodies as in the fiducial scenario, although we run simulations with all additional combinations of four impact parameters, and three impact speeds, listed in §3.2.3. The choices of $v = 0.75$ and $v = 1.5$ represent extreme scenarios to probe the sensitivity of our results to large changes in impact kinematics. All of these simulations were performed using REMIX.

Snapshots from impact simulations with different speeds and angles are shown

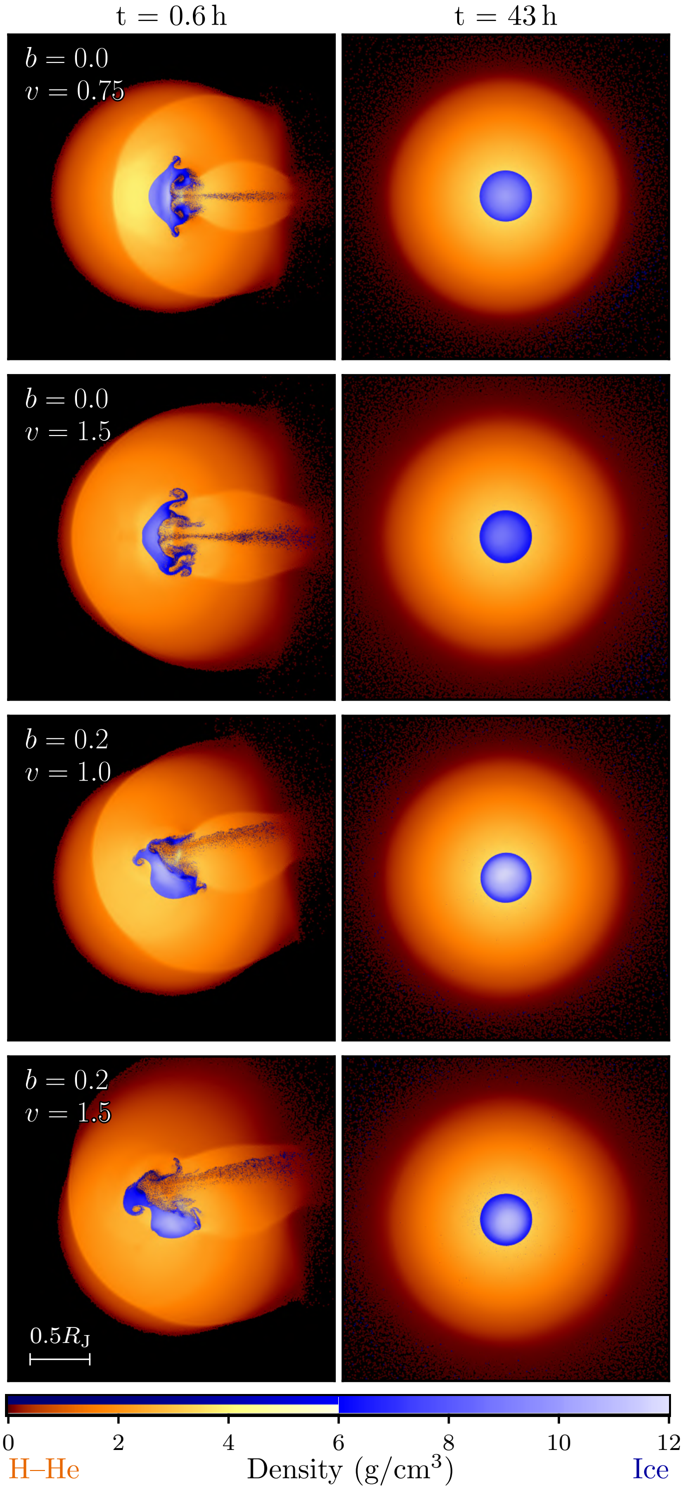

Figure 3.8: Snapshots from REMIX simulations of giant impacts onto Jupiter for different impact parameters, b, and speeds, v. Impact velocities are scaled to the mutual escape speed of the two bodies. Times correspond to core-disruption and long after impact, when heavy elements have settled to form an undiluted core. Particles are coloured by their material-type and density.

in Fig. 3.8. Although these constitute a small selection of the impacts simulated, they specifically correspond to speeds and angles in which the core is significantly disrupted. During the impacts, heavy elements mix into the envelope both through the erosion of the impactor and the disruption of the core, in particular for low impact angles. In head-on impacts energy is more effectively transferred to the core and so the core post-impact core-density in these impacts is lower than for off-axis impacts. However, for all impact configurations, heavy elements settle over short timescales to form an undiluted core with a discrete boundary to the H–He envelope.

3.4.3 Numerical resolution

Numerical resolution determines not only the minimum length scale probed in simulations, but can also significantly influence the accuracy of simulated fluid behaviour at all length scales. Total numerical convergence is not strictly achievable in simulations where turbulence at unresolved scales should grow to affect the larger system, such as in a Kelvin–Helmholtz instability across a sharp boundary. However, the convergence of large-scale simulation outcomes, rather than the specific turbulent fluid motion from which they evolved, can in many cases be demonstrated (Kegerreis et al., 2019). Not only does the computational efficiency of the SWIFT code allow us to simulate giant impacts at high resolutions, but the REMIX scheme has been demonstrated to improve numerical accuracy such that convergence can be achieved at lower resolutions than in equivalent traditional SPH simulations (§2.2).

Here, we explore numerical convergence of dilute core production – or lack thereof – in our head-on fiducial impact. We carry out simulations at resolutions of 10^n with n = 5–8.5, in steps of 0.5, SPH particles. A resolution of $10^{8.5}$ particles is, to our knowledge, the highest number of SPH particles used to model a planetary giant impact to date. All these simulations were performed using REMIX.

Mixing of heavy elements into the H–He envelope in simulations of different resolutions is shown in Fig. 3.9. We plot two times that represent the initial disruption of the core by the impactor and a point of significant mixing. The mixing of particles of different materials is observed for all resolutions. Increasing resolution allows us to capture mixing and instability growth at smaller length scales,

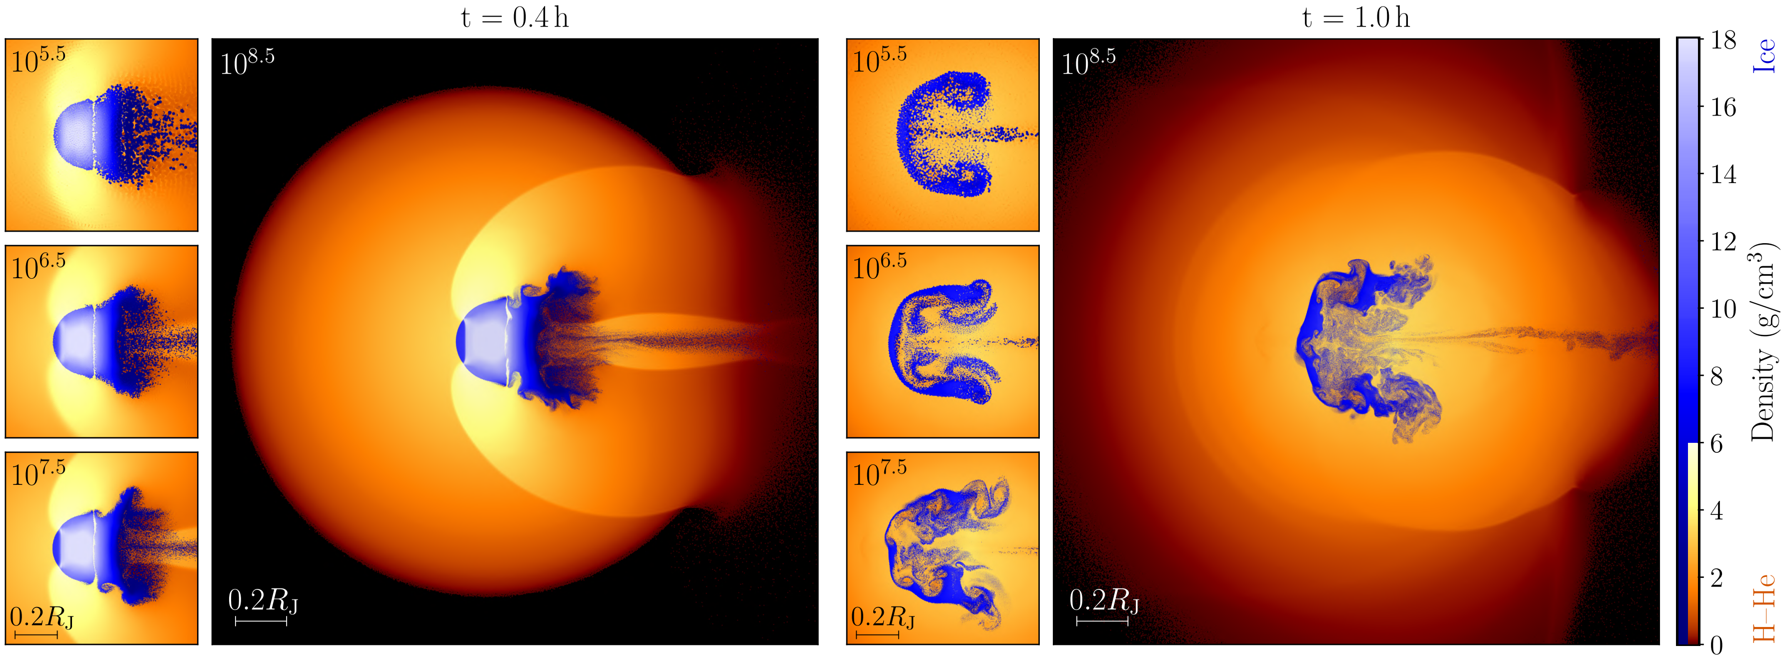

Figure 3.9: Snapshots from REMIX simulations of head-on impacts onto Jupiter, carried out with different numerical resolutions between $10^{5.5}$ and $10^{8.5}$ SPH particles. The two times plotted illustrate material mixing at different points during core-disruption. Particles are coloured by their material-type and density.

with Kelvin–Helmholtz instabilities growing at the shearing interface as predicted in §3.3.1. Because of these turbulent effects, as resolution is increased large scale features are less symmetric about the impact axis and are more significantly disrupted by fluid behaviour at smaller scales. However, it should be noted that particles in the lower resolution simulation represent larger masses, so mixing of a smaller total number of these particles can still constitute a significant degree of mixing by mass.

Due to the significant computational time required to simulate the highest resolution $10^{8.5}$, we do not simulate this impact until oscillations have fully dissipated and the planet has regained spherical symmetry. However, at the end of this simulation (21 h), heavy elements show a similar settling behaviour. In all simulations at lower resolutions than this (10^5–10^8), an undiluted and settled core is produced within the simulated time.

3.4.4 Direct Liu et al. (2019) comparison

Here we simulate impacts that are set up to directly follow the impact of L19 that they found could produce a well-mixed dilute core. By closely matching their initial conditions and the equations of state used, any remaining differences in simulation outcomes should be primarily due to the hydrodynamic methods used in the simulations. We additionally simulate an equivalent impact using more realistic EoS for both core and envelope materials. This will allow us to investigate the sensitivity of our results to the EoS used.

In the simulations of L19, planetary profiles are initially constructed using the SESAME EoS (Lyon, 1978) and then swapped for Tillotson and ideal gas EoS for the impact simulations, replacing the initial particle internal energies to recover the SESAME pressure and density profiles (S.-F. Liu, personal communication, 2020). For our subset of comparison simulations set up to match those of L19, we therefore carry out a similar process by swapping the EoS from profiles calculated using ANEOS forsterite, AQUA, and CD21 H–He to Tillotson EoS and ideal gas with $\gamma = 2$. We verify that our proto-Jupiter profile closely matches the one used in L19's simulations, so the difference between using these EoS or SESAME to construct the pre-swapped profiles is minor. We find that planets with a swapped-in ideal gas

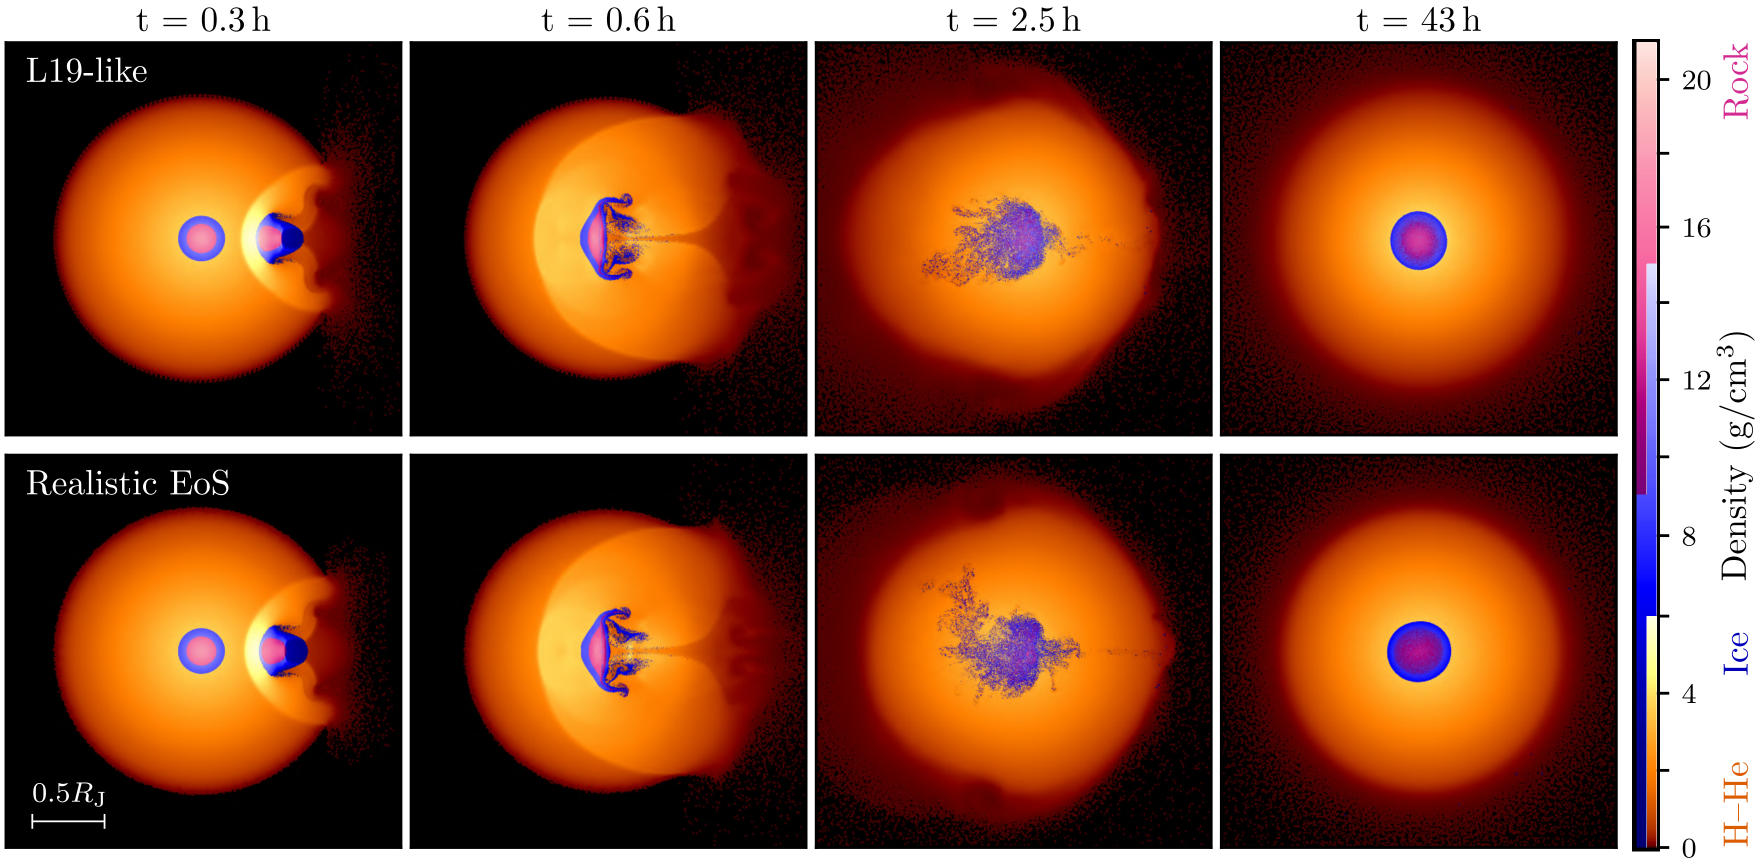

Figure 3.10: Snapshots from REMIX simulation of head-on impacts onto Jupiter set up to most closely follow the initial conditions of the impact of L19 that produced a dilute core (top) and the equivalent scenario with more realistic equations of state (bottom). Particles are coloured by their material-type and density.

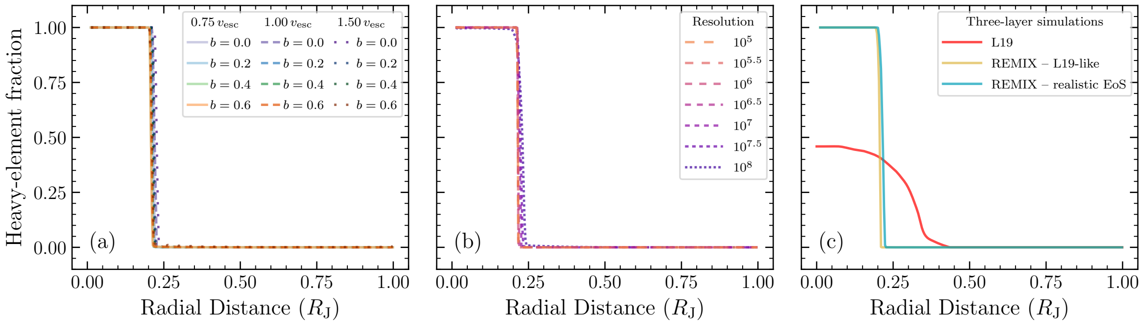

Figure 3.11: Radial profiles of heavy-element mass fraction at 43 h after impact, from simulations (a) at different impact speeds and angles; (b) with different numerical resolution; (c) set up to closely follow the initial conditions of L19, including the equivalent profile from their simulation that produced a dilute core. Heavy-element fraction is measured by the ratio of heavy-element SPH particle mass to total mass in 300 radial shells.

envelope are not stable for the duration of settling simulations. Therefore, impact simulations that use these EoS are run without prior settling simulations. As in previous simulations, the proto-Jupiter has a total mass of 308 $M_\oplus$ and the impactor has a mass of 10 $M_\oplus$.

Snapshots from impacts of both the direct L19 comparison and the equivalent simulation with improved, more sophisticated EoS are shown in Fig. 3.10. Although there are small differences in dynamics during the course of these impacts, they each follow a similar evolution, both to each other and to all previous impacts simulated here. The core is disrupted and material temporarily mixes into the envelope, but heavy elements settle to form a discrete core–envelope interface over short timescales. The rock and ice core materials remain mixed with each other at later times, however they are not diluted by envelope material.

3.5 Discussion

None of the giant impact simulations presented here produced a dilute core. The heavy-element mass fraction profiles of the post-impact planets from our simulations are plotted in Fig. 3.11. No dilute core was produced in our simulations (a) spanning a range of impact speeds and angles; (b) at different resolutions; and (c) between planets with different compositions. We carried out simulations both under conditions set up to directly mirror those of L19 and in conditions set up to facilitate mixing and remove potential barriers to it, in an attempt to offer the best chances of dilute core production. The red line in Fig. 3.11(c) shows the profile of post-impact heavy-element fraction from the simulation of L19 that produced a well-mixed dilute core with a central heavy-element fraction of $z < 0.5$. The categorical difference in our results and the simulation of L19 are likely due to differences in the simulation methodologies used.

Our simulations were largely performed using REMIX SPH, an advanced smoothed particle dynamics formulation that was developed specifically to improve the treatment of material mixing in SPH simulations. As a Lagrangian method where interpolation points move with the fluid velocity, we do not encounter the issues in regions of large bulk motion that are known to affect grid-based codes. In this aspect, the methods used by L19 are potentially limited. The well-established overmixing in grid-based codes in regions of large bulk motion through the grid may be an explanation for their results: in their simulations the core rapidly mixes into the envelope as it is accelerated by the impactor.

We have not considered mixing of material below the particle scale and the chemical processes that might affect the evolution of materials as they mix. This could be particularly relevant since core material has been predicted to be soluble in metallic hydrogen in the conditions of Jupiter's deep interior (Wilson & Militzer, 2011, 2012).

However, the impact simulation of L19 that produced a dilute core constitutes an extreme scenario with specific impact conditions. This, combined with the predictions of a dilute core in Saturn (Mankovich & Fuller, 2021) in addition to that Jupiter, might suggest that it is more likely that dilute cores are produced as part of the extended formation processes of giant planets, rather than through low-likelihood stochastic events.

3.6 Conclusions

We have presented results from REMIX SPH simulations of giant impacts onto Jupiter to investigate the feasibility of this as the process by which the planet's dilute core was formed. We varied impact speed, angle, numerical resolution, the number of layers in the pre-impact planets, and the equations of state used to represent proto-Jupiter and impactor materials. The impact dynamics in all simulations followed the same trend: initial disruption and partial mixing, followed by settling and re-formation of an undiluted heavy-element core on ∼10 h timescales. Our results contrast with the simulation of Liu et al. (2019) that produced a dilute core with a central heavy-element fraction of $Z < 0.5$ and a smooth transition to the envelope.

REMIX SPH scheme was specifically designed to improve the treatment of mixing and instability growth. Despite our approach offering optimal conditions and spanning a wide parameter space, dilute cores were not produced in any of our simulations. This result, together with observations that dilute cores are not unique to Jupiter, suggests that a single, extreme giant impact is unlikely to be the origin of dilute cores in giant planets.

CHAPTER 4

Material strength with REMIX

Up to this point, we have focused on simulating inviscid fluids. For giant impacts at the scale of Jupiter, and down as far as the scale of the impact thought to have formed the Moon, the sizes and energies involved are large enough that the treatment of solid materials as strength-less fluids is appropriate. At smaller scales, however, the solid-body strength properties of materials become increasingly significant (Holsapple & Michel, 2008). In this chapter, I present the strength models implemented in the REMIX SPH scheme to enable simulations of impact processes where material strength physics plays a crucial role.

4.1 Introduction

Continuum mechanics describes the evolution and deformation of fluids and solids, modelled as continuous media. A continuum system can be fully represented by a set of continuous fields – such as the material density, velocity, and stress – that describe the system's state, and a set of governing equations that describe its evolution.

In an astrophysical context, smoothed particle hydrodynamics (SPH) is primarily used to simulate inviscid fluids: the continuum fluid is discretised into a set of particles that evolve under the influence of gravity and hydrodynamic forces. The hydrodynamic behaviour is dictated by the Euler equations, which describe the adiabatic evolution of the inviscid fluid, driven by pressure gradients while conserving mass, momentum, and energy. For more complex materials, however, the momentum equation can be generalised to depend not only on gradients of scalar pressure, but of a tensorial description of stress, through the Cauchy momentum equation,

$$\frac{Dv^\beta}{Dt} = \frac{1}{\rho} \nabla^\alpha \sigma^{\alpha\beta} + g^\beta \ , \tag{4.1}$$

where the material derivative of velocity is calculated from the gradient of the stress tensor, σ, the density, ρ, and accelerations due to body forces, $\mathbf{g}$, that could for instance include gravity. Note that for SPH, where particles move with the velocity of material at their positions, material derivatives directly correspond to the derivatives of fields sampled at moving particle positions. The stress tensor captures responses to deformation as well as to compression. The system's more complex governing physics is incorporated in the stress tensor, allowing the description of, for instance, viscous Newtonian fluids, non-Newtonian fluids, and solid bodies.

The influence of material strength physics on planetary impact phenomena becomes more significant at smaller scales of size and energy. For these simulations, the SPH approach offers additional advantages over grid-based methods, beyond those previously discussed for fluids, because particles inherently track the material's history[1]. The SPH framework can be adapted to model material strength physics to simulate solid-state mechanics by using an elastic–perfectly-plastic constitutive model (Libersky & Petschek, 1991); more sophisticated models for plastic deformation (Collins et al., 2004); models that consider the growth of fractures on the sub-particle scale (Benz & Asphaug, 1994, 1995; Jutzi, 2015); and models for porosity (Jutzi et al., 2008). These methods have been applied in simulations of planetary impacts to simulate impacts across a wide range of scales (Asphaug & Benz, 1996; Ballantyne et al., 2023; Emsenhuber et al., 2018; Raducan & Jutzi,

---

[1] For instance, the method described in §4.2.6 by which fracture growth is seeded by flaws and associated activation thresholds, cannot easily be adapted to non-Lagrangian codes (Collins et al., 2004).

2022). However, grid-based codes, such as the finite-difference code iSALE (Collins et al., 2004), are more frequently utilised for smaller scale cratering impacts where the effects of self-gravity are negligible.

Here, we describe the material strength models implemented in the REMIX scheme (§4.2) and apply them to some test simulations to assess the performance of REMIX in these scenarios and to investigate the effect of different physical and numerical models on simulation outcomes (§4.3).

4.2 Strength models

Here we present the models of material strength that we use to enable the treatment of solid-state mechanics with REMIX SPH. We describe the physical motivation of each model and their implementation into the REMIX equations.

4.2.1 Equations of motion

To simulate solid-body mechanics, the REMIX equations of motion must be adapted from depending only on gradients of the isotropic, scalar pressure, to those of a tensorial representation of stress. The stress tensor of an SPH particle i, σ_i, is given by

$$\sigma_i^{\alpha\beta} = -P_i \delta^{\alpha\beta} + S_i^{\alpha\beta} \ , \tag{4.2}$$

where the hydrodynamic pressure, P, relates to changes in volume and the deviatoric stress tensor, S, describes the material's response to changes in shape. The deviatoric stress tensor is traceless and so the hydrostatic pressure can be expressed in terms of invariants of the stress tensor by $P = -\mathrm{tr}(\sigma)/3$ in three dimensions, where the trace $\mathrm{tr}(\sigma) \equiv \sigma^{\alpha\alpha}$.

The REMIX equations of motion, modified to incorporate the stress tensor, are expressed as

$$\frac{d\rho_i}{dt} = \sum_j m_j \frac{\rho_i}{\rho_j} v_{ij}^{\alpha} \frac{1}{2}\left(\left.\frac{d\widetilde{\mathcal{W}}}{dr^{\,\alpha}}\right|_{ij} - \left.\frac{d\widetilde{\mathcal{W}}}{dr^{\,\alpha}}\right|_{ji} \right) + \left(\frac{d\rho_i}{dt}\right)_{\rm difn} + \left(\frac{d\rho_i}{dt}\right)_{\rm norm} , \tag{4.3}$$

$$\frac{dv_i^{\alpha}}{dt} = \sum_j m_j \frac{\sigma_{ij}^{\alpha\beta} - Q_{ij}\delta^{\alpha\beta} + \sigma_{ji}^{\alpha\beta} - Q_{ji}\delta^{\alpha\beta}}{\rho_i\,\rho_j} \frac{1}{2}\left(\left.\frac{d\widetilde{\mathcal{W}}}{dr^{\,\beta}}\right|_{ij} - \left.\frac{d\widetilde{\mathcal{W}}}{dr^{\,\beta}}\right|_{ji} \right) , \tag{4.4}$$

$$\frac{du_i}{dt} = -\sum_j m_j \frac{\sigma_{ij}^{\alpha\beta} - Q_{ij}\delta^{\alpha\beta}}{\rho_i\,\rho_j} v_{ij}^{\alpha} \frac{1}{2}\left(\left.\frac{d\widetilde{\mathcal{W}}}{dr^{\,\beta}}\right|_{ij} - \left.\frac{d\widetilde{\mathcal{W}}}{dr^{\,\beta}}\right|_{ji} \right) + \left(\frac{du_i}{dt}\right)_{\rm difn} , \tag{4.5}$$

which return to the purely hydrodynamic REMIX equations of motion for $\sigma^{\alpha\beta} = -P\delta^{\alpha\beta}$, corresponding to $S^{\alpha\beta} = 0$, for all particles. Here the stress tensors are expressed as pairwise quantities, σ_{ij} and σ_{ji}, since as we incorporate physical models these describe different interactions depending on the phases of particles in each pair (§4.2.2) and require pairwise artificial corrections to deal with numerical tensile instabilities (§4.2.4).

When simulating solid bodies with SPH, using an evolved density estimate rather than an integral form is useful for avoiding kernel smoothing of free surfaces (Benz & Asphaug, 1994; Gray et al., 2001; Schäfer et al., 2016). The REMIX formulation already uses an evolved density estimate and so is well set up for the incorporation of material strength physics, all of the models for which are contained within the stress tensor in the equations above. Additionally, for solid particles, we do not include the normalising term in the density evolution, since we find that this treatment leads to artificial cohesive effects in elastic simulations with multiple bodies, such as between the colliding cylinders presented in §4.3.1.

In later sections, we include comparison simulations that were carried out using a traditional SPH formulation ("tSPH"). This formulation is based on the equations presented in Appendix C, however instead of using the integral form of the density estimate, the density is evolved by the equation,

$$\frac{d\rho_i}{dt} = \sum_j m_j\, v_{ij}^{\alpha} \left.\frac{\partial W}{\partial r^{\alpha}}\right|_{ij} , \tag{4.6}$$

which is thermodynamically consistent with the tSPH energy evolution equation. Similarly to the REMIX equations of motion above, pressures have been replaced by stress tensors, which are calculated exactly as those that are used in equivalent REMIX simulations.

4.2.2 Interactions of fluids and solids

The equations of motion presented in the previous section combine the hydrodynamics treatment of REMIX with solid-body dynamics, and can therefore treat both fluids and solid material. In our implementation, the models for material strength outlined in the following sections are only applied to particles that are currently in a solid state. Additionally, solid materials are able to sustain negative pressures – lower than the vacuum pressure – when under tension. Although liquids can occupy metastable states with negative pressure, these states are unstable to cavitation, which leads to the formation of vapour bubbles and ultimately raises the pressure back above zero (Caupin & Herbert, 2006). Since we do not treat cavitation on the particle scale, the pressure of fluid particles are always greater than or equal to zero. Additionally, the deviatoric stress of fluid particles is fixed to zero, since we still treat fluids as inviscid.

With the relatively simplistic thermodynamic treatment used in the test simulations presented here – using Tillotson equations of state (Melosh, 1989) and no latent heats – particles switch phase between fluid and solid as their temperature transitions through the constant melting temperature.

For conservation of momentum and energy, the momentum equation (Eqn. 4.4) must be antisymmetric in the exchange of i and j for each particle. Therefore, without a specific treatment, the interaction of a fluid–solid particle pair leads to behaviour that is halfway between that of a fluid and a solid. This is not physically appropriate, as it would lead to effects like the build up of stress as a fluid flows across a solid surface, as well as sustained tensile interaction between fluids and solids.

To address these interactions between particles of different phases, we treat fluid–solid interactions as if both particles were of a strengthless fluid. Therefore, the pairwise stress tensor used in Eqns. 4.3–4.5 is given by

$$\sigma_{ij}^{\alpha\beta} = \begin{cases} \sigma_i^{\alpha\beta} & \text{for solid } i \text{ and } j \ , \\ -P_i\delta^{\alpha\beta} & \text{if } P_i > 0, \text{ for fluid } i \text{ and/or } j \ , \\ 0 & \text{if } P_i \leq 0, \text{ for fluid } i \text{ and/or } j \ . \end{cases} \tag{4.7}$$

Additionally, the evolution of the deviatoric stress (described in the following sections) only considers interactions between both-solid particles, which means that fluid flowing over a solid surface will not accumulate stress in the same way as a distorting solid body.

4.2.3 Elastic model

Under perfectly elastic behaviour, a solid body will accumulate stress as it deforms and will evolve to attempt to regain its original, unstressed shape. An elastic model, based on the continuum Hooke's law in which the accumulated stress is proportional to the distortion (or strain) of the material, forms the basis of all material strength models presented here (Libersky & Petschek, 1991).

The rate of deformation of a continuous medium is described by the strain-rate tensor, $\dot{\epsilon}$, which we calculate at the position of particle i, by

$$\dot{\epsilon}_i^{\alpha\beta} = \frac{1}{2}\left(\left.\frac{dv^\alpha}{dr^\beta}\right|_i + \left.\frac{dv^\beta}{dr^\alpha}\right|_i \right) \ , \tag{4.8}$$

where velocity gradients are calculated using the REMIX kernel gradient terms,

$$\left.\frac{dv^\alpha}{dr^\beta}\right|_i = \sum_j \phi_{ij} \left(v_j^\alpha - v_i^\alpha\right) \left.\frac{d\widetilde{\mathcal{W}}}{dr^\beta}\right|_{ij} \frac{m_j}{\rho_j} \ . \tag{4.9}$$

The phase parameter ϕ_{ij} takes the value $\phi_{ij} = 1$ if both i and j are solid, and 0 otherwise. Therefore only the deformation of solid material is considered.

In the absence of rotation, the perfectly elastic evolution of the deviatoric stress tensor can be calculated from only the material properties and the strain-rate tensor. However, in general, the material may experience local rotation that does not act to deform it, but in the fixed coordinate system of the simulation, contributes to changes in the deviatoric stress tensor.

For example, consider a rubber band that is stretched in the x-direction. If this rubber band is rotated as a rigid body by 90° such that it is now extended in the y-direction, without additional deformation, then releasing the band will cause it to snap back in the y-direction to regain its original shape. Therefore, in a fixed coordinate system, the stress tensor must evolve during the rotation; otherwise the rubber band's response would be in the x-direction. In a coordinate system that rotates along with the rubber band, however, the stress tensor would not change during the rotation.

In general, every position in a continuous material can experience varying local deformation and rotation. It is the local rotation at a given position that contributes to the evolution of the deviatoric stress at that position. The local rotation rate at the position of a particle is given by the tensor

$$R_i^{\alpha\beta} = \frac{1}{2}\left(\left.\frac{dv^\alpha}{dr^\beta}\right|_i - \left.\frac{dv^\beta}{dr^\alpha}\right|_i\right) . \tag{4.10}$$

The elastic evolution of deviatoric stress can then be shown to take the form

$$\frac{dS_i^{\alpha\beta}}{dt} = 2\mu\left(\dot{\epsilon}_i^{\alpha\beta} - \frac{1}{3}\delta^{\alpha\beta}\dot{\epsilon}_i^{\gamma\gamma}\right) + S_i^{\alpha\gamma}R_i^{\gamma\beta} - R_i^{\alpha\gamma}S_i^{\gamma\beta} , \tag{4.11}$$

where μ is the shear modulus that describes the elastic stiffness of the material, and the rotation terms, which depend on R_i, factor in the local rotation at the position of the particle (Dienes, 1979).

4.2.4 Artificial stress

Solid materials can experience tensile stress. In the SPH representation, negative pressures in material that is under tension can lead to the spurious clustering of SPH particles. This can lead to numerical fragmentation and artificial growth of fractures, due to a tensile instability, that can significantly affect the treatment of simulated materials.

This tensile instability can by addressed by introducing "artificial stress" that acts as a short-range repulsive force to prevent the clumping of material in tension. Monaghan (2000) introduced an approach that adds an additional term to the mo-

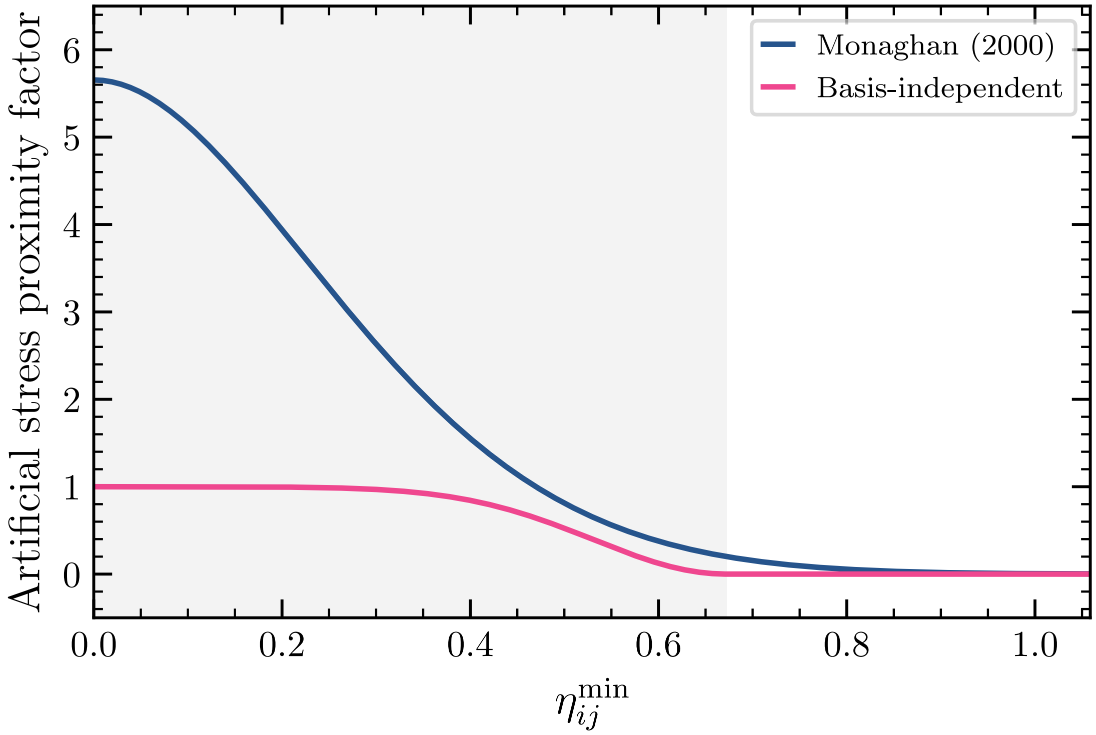

Figure 4.1: The functions that scale the strength of artificial strength, plotted against particle separation scaled to the larger of the two smoothing lengths. For the Monaghan (2000) approach, the factor plotted is $0.2\, f_{ij}^4$, with f_{ij} given by Eqn. 4.13. For the basis-independent approach that detects proximity similarly to the artificial viscosity and diffusion schemes of REMIX, the factor plotted is f_{ij} (Eqn. 4.16). A Wendland C^2 kernel is used with $\eta_{\rm kernel} = 1.487$. The shaded region corresponds to the region $\eta_{ij}^{\rm min} \equiv \Delta p \,/\, h_{ij}^{\rm max} < 1/\eta_{\rm kernel}$

mentum equation that modifies positive elements of the stress tensor with increasing strength as particles approach each other. This approach is relatively simple and is frequently adopted by SPH schemes that simulate material strength (Pearl et al., 2022; Schäfer et al., 2016). However, this approach is not basis-independent and therefore the correction applied depends upon the choice of simulation coordinate system. Gray et al. (2001) improved on this approach by applying the correction to the eigenvalues of the stress tensor, or principal stresses, rather than to individual elements of the tensor in the basis of the simulation coordinate system. However, the computational implementation of this method is more complex, because it requires the calculation of eigenvectors that are needed to transform back to the basis of the simulation coordinate system, as well as the principal stresses.

Here we describe the implementation of both the Monaghan (2000) artificial stress into REMIX and present a novel approach that, like the method of Gray et al. (2001), is based on the modification of principal stresses although does not require the additional steps of eigenvector calculation and basis transformation.

As well as modifying the momentum equation, as done in past approaches, for both these methods we also modify the energy evolution equation to make these methods conservative in our implementation. For conservation, the artificial stress appears in the equations of motion as a pairwise modification to the stress tensor, $\tilde{\sigma}_{ij} = \sigma_i + \sigma_{\mathrm{art},\,ij}$, which is used in place of σ_i in Eqn. 4.7, for solid–solid interactions. With this implementation, the momentum equation only differs from the Monaghan (2000) terms by notation (in the grouping of densities) and in the densities used in these terms with REMIX, based on the different forms of the equations of motion.

These methods are tested with both REMIX and tSPH in simulations of colliding cylinders in §4.3.1. Since internal energy is not used in the equation of state in that particular test, which is often used to assess the approach to dealing with the tensile instability, this modification to make the equations of motion conservative does not differentiate our Monaghan (2000) implementation from similar simulations of past work. However, we do not use XSPH, which is often combined with artificial stress as a correction for these specific simulations (Gray et al., 2001; Monaghan, 2000; Pearl et al., 2022; Schäfer et al., 2016). With XSPH, particles are moved with a velocity different from the sampled velocity of the material, meaning it is not a Lagrangian method. The velocity used to move the particles is based on a kernel-smoothed estimate of the material velocity. This approach helps address the tensile instability, as individual particles are less likely to stray far from the bulk motion at the kernel scale.

Monaghan (2000) artificial stress

For the Monaghan (2000) approach, the artificial stress is applied to positive elements of the stress tensor individually, by

$$\sigma^{\alpha\beta}_{\text{art},\,ij} = \begin{cases} -\epsilon_s\, f^n_{ij}\, \sigma^{\alpha\beta}_i & \text{if } \sigma^{\alpha\beta}_i > 0\ , \\ 0 & \text{otherwise}\ , \end{cases} \tag{4.12}$$

where $\epsilon_s = 0.2$ and $n = 4$ are constants and f_{ij} is a function that increases as particles approach each other, given by

$$f_{ij} = \frac{W\big(r_{ij},\ h^{\max}_{ij}\big)}{W\big(\Delta p,\ h^{\max}_{ij}\big)}\ , \tag{4.13}$$

where $r_{ij} \equiv |\mathbf{r}_i - \mathbf{r}_j|$, $h^{\max}_{ij}$ is the larger of h_i and h_j, and Δp is a measure of the expected separation of the closest neighbour, for which we use $\Delta p = h^{\max}_{ij} / \eta_{\text{kernel}}$, similarly to the η_{crit} used in the artificial viscosity and diffusion schemes of REMIX. Since $f^n_{ij} \sim W\left(r_{ij},\ h^{\max}_{ij}\right)^4$, this function peaks sharply as the particle separation is reduced. The function $0.2\, f^4_{ij}$ is plotted in Fig. 4.1 as a function of $\eta^{\min}_{ij} \equiv r_{ij} \,/\, h^{\max}_{ij}$ for a Wendland C^2 kernel with $\eta_{\text{kernel}} = 1.487$. This function takes a value 5.7 for $r_{ij} = 0$, and a value 0.2 for $r_{ij} = \Delta p$ by the choice of ϵ_s.

With this form, the artificial stress applied depends on the basis in which the condition $\sigma^{\alpha\beta}_i > 0$ is examined, since this element-wise, conditional operation is not basis-independent. This can be illustrated by a simple example. Let us consider a matrix that here plays the role of an unmodified stress tensor,

$$M = \begin{pmatrix} -2 & 1 & 0 \\ 1 & -2 & 0 \\ 0 & 0 & -1 \end{pmatrix}\ , \tag{4.14}$$

which has eigenvalues $\lambda_1 = -3$, $\lambda_2 = -1$, $\lambda_3 = -1$. Under an operation like that described by Eqn. 4.12, a correction would be applied to the positive, off-diagonal elements of M. However, when rotated to the eigenbasis, there are no positive elements – this choice of basis transformation is for simplicity, since it is clear that all eigenvalues are negative. Therefore, in one basis a correction would be applied and in another no correction would be applied. From this simple example we see that the artificial stress applied between each particle pair depends on the coordinate system that is chosen as the basis for these calculations.

A new basis-independent artificial stress approach

Here we present a new artificial stress method that is basis-independent and that aims to apply corrections more closely targeted at problematic particles. Like the approach of Gray et al. (2001), we apply a correction based on the eigenvalues of the stress tensor. However, instead of treating each principal stress separately, we modify all three principal stresses by the same value based on the largest of the eigenvalues, $\sigma_i^{\max}$. The added correction is therefore proportional to $\sigma_i^{\max}\,\delta^{\alpha\beta}$ in an arbitrary basis, since adding a scalar multiple of the Kronecker delta is a basis-independent operation. Therefore, this correction can be applied in the basis of the coordinate system, without the need for basis transformations and the calculation of eigenvectors that would otherwise be needed for the transformation.

With this approach, the artificial stress takes the form

$$\sigma^{\alpha\beta}_{\mathrm{art},\,ij} = \begin{cases} -f_{ij}\,\sigma_i^{\max}\,\delta^{\alpha\beta} & \text{if } \sigma_i^{\max} > 0\ , \\ 0 & \text{otherwise}\ , \end{cases} \tag{4.15}$$

where we additionally use a new form for f_{ij} to switch on the strength of the artificial stress as particles approach each other, which is not directly tied to the shape of the choice of kernel function and which peaks at a value of 1 for a more controlled correction, given by

$$f_{ij} = \begin{cases} 1 - \exp\left[-\left(\dfrac{\eta_{ij}^{\min} - \eta_{\mathrm{crit}}}{0.2}\right)^2\right] & \text{for } \eta_{ij}^{\min} < \eta_{\mathrm{crit}}\ , \\ 0 & \text{otherwise}\ , \end{cases} \tag{4.16}$$

with $\eta_{\mathrm{crit}} = 1/\eta_{\mathrm{kernel}}$. This form is based on the function used to switch off the linear reconstruction of velocities, internal energies and densities in the artificial viscosity and diffusion schemes in REMIX. Unlike the function used for this purpose in the Monaghan (2000) method above, this function takes a value of 0 at the expected position of the nearest neighbour and only reaches a value of 1 for $r_{ij} = 0$, as plotted in Fig. 4.1. Therefore, the correction is more targeted towards particles that are problematic and the magnitude of the correction itself is constrained by the maximum, positive principal stress. A comparison of these two artificial stress approaches in simulations of colliding cylinders is presented in §4.3.1.

4.2.5 Plastic deformation

Up to this point, we have considered perfectly elastic materials and the challenges faced in modelling solids in SPH. However, a real solid will not behave elastically under all conditions: under enough stress, a solid body will permanently deform and not regain its original shape when allowed to relax. To model this plastic deformation, we consider two previously tested approaches to modelling plastic deformation based on a yield stress, Y, which characterises the transition from elastic to plastic behaviour, and the yield criterion, by which Y is applied. In addition to these models, we describe methods for softening the material as temperature approaches the melting temperature and for low densities.

With both approaches described here, the second invariant of the deviatoric stress,

$$J_2 = \frac{1}{2} S^{\alpha\beta} S^{\beta\alpha} \, , \tag{4.17}$$

is used as a measure of the magnitude of shear stress experienced by the material. J_2 is compared with the yield stress to determine whether material behaves elastically or plastically and the deviatoric stress is multiplied by a factor f_Y to capture this behaviour, by

$$S_Y^{\alpha\beta} = f_Y S^{\alpha\beta} \, , \tag{4.18}$$

where $0 \leq f_Y \leq 1$. When these models for plastic deformation are used, $S_Y^{\alpha\beta}$ is calculated each time the deviatoric stess is evolved in time. $S_Y^{\alpha\beta}$ replaces $S^{\alpha\beta}$ for all calculations, and acts to reduce the deviatoric stress as the material yields. Additionally, we limit the maximum negative pressure by $P_{\rm min} = -Y_0$, where Y_0 is the yield stress at zero pressure (Luther et al., 2022).

A constant yield stress approach

We first consider a simple approach with a constant yield stress, $Y = Y_0$, and the von Mises yield criterion (Benz & Asphaug, 1994). The deviatoric stress is modified by

$$f_Y = \begin{cases} \dfrac{Y_0^2}{3J_2} & \text{for } 3J_2 > Y_0^2 \ , \\ 1 & \text{otherwise} \ . \end{cases} \tag{4.19}$$

With this approach the yield stress of the material does not depend on pressure or the damaged state of the material.

A pressure-dependent yield stress approach

Next, we consider the pressure dependent yield stress of Collins et al. (2004). This approach also models changes in the yield stress, as a material is allowed to transition from a cohesive body to a fragmented state as it gets damaged. This fracturing is modelled by the parameter D which evolves from 0 to 1 as a material is damaged. The treatment of the accumulation of damage for an SPH particle, D_i, is described later in §4.2.6.

With this approach, intact material has a yield stress

$$Y_{in} = Y_0 + \frac{\mu_i P}{1 + \dfrac{\mu_i P}{Y_M - Y_0}} \ , \tag{4.20}$$

where Y_0 is the yield stress at zero pressure, μ_i is the coefficient of internal friction of the material, and Y_M is the von Mises plastic limit of the material: $Y_{in} \to Y_M$ as $P \to \infty$ (Lundborg, 1968).

The yield stress of damaged material follows a Coulomb dry-friction law, limited by the yield stress of the intact material, via

$$Y_d = \begin{cases} \mu_d P & \text{for } \mu_d P < Y_{in} \ , \\ Y_{in} & \text{otherwise} \ , \end{cases} \ , \tag{4.21}$$

where μ_d is the coefficient of friction of the damaged material.

The yield stress of intact and damaged materials combine for intermediate values of D according to

$$Y = (1 - D)\,Y_{in} + D\,Y_d\ . \tag{4.22}$$

This yield stress is applied to the deviatoric stress (by Eqn. 4.18) with the factor

$$f_Y = \begin{cases} \dfrac{Y}{\sqrt{J_2}} & \text{for } \sqrt{J_2} > Y\ , \\ 1 & \text{otherwise}\ , \end{cases} \tag{4.23}$$

In §4.3.2 we test the effect of using either Eqn. 4.19 or Eqn. 4.23 as a yield criterion, for a fixed yield stress. Future work will involve testing this approach in crater-forming simulations into brittle material, to validate both the pressure- and damage-dependent aspects of this method.

Thermal weakening

The shear strength of a material weakens as the temperature approaches the melting temperature. To capture this temperature dependence, we multiply the yield stress of the material by a factor that switches to zero at the melting temperature, by

$$Y_T = \begin{cases} Y \tanh\left[\xi\left(\dfrac{T_m}{T} - 1\right)\right]\ , & \text{for } T < T_m\ , \\ 0 & \text{otherwise}\ , \end{cases} \tag{4.24}$$

where T is the temperature of the material and T_m is the melting temperature, and ξ is a material constant (Collins et al., 2004; Emsenhuber et al., 2018; Ohnaka, 1995). We use $\xi = 1.2$ for the simulations presented here (Emsenhuber et al., 2018; Luther et al., 2022). Y_T replaces Y when used in either Eqn. 4.19 or 4.23 depending on which method is used.

Density weakening

Additionally, to capture the weakening of low-density, solid material, we modify the yield stress by

$$Y_\rho = \begin{cases} Y\left(\dfrac{\rho}{\rho_\text{weak}}\right)^4 & \text{for } \rho < \rho_\text{weak} \ , \\ Y & \text{otherwise} \ , \end{cases} \tag{4.25}$$

where we take $\rho_\text{weak} = 0.85\rho_0$ where ρ_0 is the reference density of the material (Luther et al., 2022). Y_ρ replaces Y in either Eqn. 4.19 or 4.23 depending on which method is used, and can be combined with the thermal weakening, so that the yield stress is reduced both at high temperatures and at low densities.

4.2.6 Tensile fracture model

Here we describe the method we use to model how fractures grow in material under tension (Benz & Asphaug, 1994, 1995). The state of damage of a material is tracked by a parameter D that takes values between $D = 0$, corresponding to an intact solid, to $D = 1$, which corresponds to maximally damaged material (Grady & Kipp, 1980). A benefit of the Lagrangian SPH approach is that the evolution history of each part of the material is inherently tracked. Damage accumulates as tensile stresses activate the growth of fractures within the material and directly affects the elements of the stress tensor, and can indirectly affect them depending on the other models used, for example through the yield stress as described in §4.2.5.

For this model, each SPH particle i is initialised with a number of flaws, N_i^flaw, that are associated with the particle for the duration of the simulation and are not resolved on the particle scale. Each flaw, n, has an associated failure strain, ϵ_n^act. If the SPH particle experiences strains above ϵ_n^act, the flaw n will be activated and fractures will grow. Flaws are distributed in a solid, consisting of many SPH particles, by a Weibull distribution (Weibull, 1939): particles in the solid are selected at random and assigned a flaw until each particle has at least one flaw, giving a total number of flaws in the solid $N_\text{solid}^\text{flaw}$. The activation threshold of each flaw $n = 1, 2, \ldots, N_\text{solid}^\text{flaw}$, the associated failure strain, is given by

$$\epsilon_n^\text{act} = \left(\frac{n}{k_W V}\right)^{1/m_W} \ , \tag{4.26}$$

where V is the total volume of the brittle solid, and k_W and m_W are material-dependent constants for the Weibull distribution.

As the material evolves in the simulation, SPH particles will experience tensile stresses that will activate the seeded flaws, leading to the localised accumulation of damage. The local scalar strain is calculated from the maximum principal stress, $\sigma_i^{\max}$, by

$$\epsilon_i = \begin{cases} \dfrac{\sigma_i^{\max}}{(1 - D_i)\, E} & \text{if } \sigma_i^{\max} > 0 \,, \\ 0 & \text{otherwise} \,, \end{cases} \tag{4.27}$$

where D_i is the current damage accumulated by the particle and E is the material's Young's modulus, which can be expressed in terms of the shear modulus, μ, and the bulk modulus, K, by $E = 9K\mu/(3K + \mu)$. While $\epsilon_i > \epsilon_n^{\mathrm{act}}$ for the particle i containing the flaw n, then the flaw n is activated. If ϵ_i drops below $\epsilon_n^{\mathrm{act}}$, then the flaw is not currently activated.

Damage accumulates based on the growth of individual cracks within a sphere of radius R_s, growing with speed c_g. For each particle, the number of cracks growing is based on the number of currently activated flaws for that particle, N_i^{act}. The accumulation of damage can then be expressed by

$$\frac{dD_i^{1/3}}{dt} = N_i^{\mathrm{act}} \frac{c_g}{R_s} \,. \tag{4.28}$$

Note that here the damage of each particle can only increase as it becomes more damaged with time. We use $R_s = \sqrt[3]{m_i/\rho_i}$ and the speed of crack growth is taken to be 0.4 times the velocity of a longitudinal elastic wave, and is therefore given by

$$c_g = 0.4\sqrt{\frac{K + \frac{4}{3}(1 - D_i)\, \mu}{\rho_i}} \,, \tag{4.29}$$

calculated using the shear modulus and bulk modulus, as well as the density and damage of the particle (Benz & Asphaug, 1995).

However, at any given time, the growth of damage is capped by the number of active flaws,

$$D_{\max,\,i} = \frac{N_i^{\text{act}}}{N_i^{\text{flaw}}}\ , \tag{4.30}$$

so that the accumulation of damage is not dominated by the flaw with the lowest activation threshold (Schäfer et al., 2016). At a given time, damage can only accumulate up to the current $D_{\max,\,i}$. D_i is not affected if $D_{\max,\,i}$ drops below the current D_i, other than that it will not continue to grow until more flaws are activated and $D_{\max,\,i}$ increases above D_i.

For the methods tested here, the damage is applied to the stress tensor in different ways based on the models used. With the constant yield stress approach in §4.2.5, fully-damaged material behaves as a fluid, and so the stress tensor takes the form

$$\sigma_{i,\,D} = \begin{cases} -P_i\,\delta^{\alpha\beta} + (1-D_i)\ S_i^{\alpha\beta} & \text{if } P_i \geq 0\ , \\ -\,(1-D_i)\ P_i\,\delta^{\alpha\beta} + (1-D_i)\ S_i^{\alpha\beta} & \text{if } P_i < 0\ . \end{cases} \tag{4.31}$$

Here both the deviatoric stress tensor and negative pressures are modified by a factor of $(1-D_i)$. With the Collins et al. (2004) pressure-dependent yield stress approach (§4.2.5), the reduction of the shear strength of the material as it fractures is already incorporated in the yield stress (Eqn. 4.22). Therefore, damage is only applied directly to the stress tensor as a factor multiplying negative pressures

$$\sigma_{i,\,D} = \begin{cases} -P_i\,\delta^{\alpha\beta} + S_i^{\alpha\beta} & \text{if } P_i \geq 0\ , \\ -\,(1-D_i)\ P_i\,\delta^{\alpha\beta} + S_i^{\alpha\beta} & \text{if } P_i < 0\ . \end{cases} \tag{4.32}$$

The damaged stress tensor $\sigma_{i,\,D}$ is used in place of σ_i for all calculations.

4.3 Strength tests

In this section we validate the REMIX implementation of the material strength models described above. Each test scenario is chosen to investigate a specific aspect of the models and the methods introduced for each test build on top of the methods that are validated previously.

We present results from the following test scenarios:

- collisions between rubber cylinders (§4.3.1), to test the treatment of elasticity with REMIX and tSPH, including an investigation into artificial stress models;
- crater-forming impacts into aluminium (§4.3.2), to test the treatment of plasticity in impacts into strong, ductile material;
- impacts into a basalt sphere (§4.3.3), to test the treatment of fracturing in impacts into brittle material.

In each case, we investigate the effect of the model used, where multiple alternatives have been described in the previous section, as well as the effect of subtle considerations in how these models are applied.

As with the hydrodynamic tests presented in Chapter 2, these tests are all carried out in 3D, to validate our methods for 3D science applications, and use a Wendland C^2 kernel with $\eta_{\text{kernel}} = 1.487$. For aluminium and basalt sphere impacts, we use Tillotson equations of state with the parameters listed in Table 4.1.

4.3.1 Colliding rubber cylinders

We use simulations of colliding rubber cylinders to test the implementation of the elastic model, described in §4.2.3, and the artificial stress models, described in §4.2.4, in simulations using both REMIX and tSPH. This test has been adopted, in particular, to investigate whether the simulation methods are able to deal with the tensile instability that leads to the growth of artificial fractures in SPH simulations when particles can be in tension (Monaghan, 2000). To combat this, past SPH methods use artificial stress, as we also do here, but also require the "XSPH" treatment to avoid artificial fragmentation (Monaghan, 1989).

With XSPH, particles are moved with a velocity based on an average of the velocities of neighbouring particles, rather than the fluid velocity at the particle location. Additionally, when XSPH is used for validation of SPH methods with strength, it is typically used only for simulations like the colliding rubber cylinders and turned off for further validation in more general scenarios. Therefore, validation of the elastic model for colliding rubber cylinders using XSPH is not representative of the typical numerical methods used for most applications, where XSPH is not

Table 4.1: Tillotson equation of state parameters and material strength parameters. Strength parameters are given only for the models used in the relevant test simulations. Pressure is limited by $-Y_0$ and is decreased to zero for low densities by a linear function that switches from 1 to 0 between densities 0.925 ρ_0 and 0.875 ρ_0. The lucite Tillotson EoS parameters are those used by Schäfer et al. (2016) and differ from the basalt parameters only in ρ_0.

	Aluminium	Basalt	Lucite
ρ_0 (kg m^{-3})	2700	2700	1180
a	0.5	0.5	0.5
b	1.63	1.50	1.50
A (GPa)	75.2	26.7	26.7
B (GPa)	65.0	26.7	26.7
u_0 (MJ)	5.0	487	487
$u_{\rm iv}$ (MJ)	3.0	4.72	4.72
$u_{\rm cv}$ (MJ)	13.9	18.2	18.2
α	5	5	5
β	5	5	5
C_V (J kg^{-1} K^{-1})	900	790	790
T_m (K)	933	-	-
μ (GPa)	24.0	22.7	0.073
K (GPa)	-	26.7	-
Y_0 (GPa)	0.2	3.5	0.01
k_W (cm^{-3})	-	5×10^{28}	-
m_W	-	8, 8.5	-

used. Therefore, here we carry out these tests without using XSPH.

We follow the initial conditions used by Monaghan (2000) (hereafter M00), but set up our cylinders in a 3D periodic box with a small, periodic depth dimension, instead of in 2D, to validate our methods for 3D science simulations. We set up particles in two cylinders with inner radius 3 cm and outer radius 4 cm. Particles are arranged in a cubic lattice with 0.1 cm spacing. The cylinders are 15 particles deep, spanning the depth of the periodic box. Pressures are calculated by

$$P = c_0^2 \; (\rho - \rho_0) \;, \tag{4.33}$$

with sound speed $c_0 = 852$ m s^{-1} and reference density $\rho_0 = 1$. The particles are initialised with $\rho = 1$ and so are in pressure equilibrium with the surrounding region with no particles, and the cylinders are initialised with relative speed $0.12\, c_0$. A shear modulus of $\mu = 0.22\, \rho_0 c_0^2$ is used for the evolution of deviatoric stress by

Eqn. 4.11.

The collision of rubber cylinders when simulated using REMIX is shown in Fig. 4.2. Rows show the collision when simulated with the basis-independent approach to artificial stress (§4.2.4), the M00 artificial stress formulation (§4.2.4), and with no artificial stress. Without artificial stress we see spurious fractures form as the SPH particles clump when under tension. These fractures lead to the shattering of the cylinders and the evolution is not elastic, as it should be based on the physical models used. With both approaches to artificial stress, the cylinders are able to flex and sustain tension without fracturing. The desired elastic collision is observed in both cases. However, we note that with the M00 approach, here we had to set a maximum smoothing length of $h_{\text{max}} = 2.2$ cm to avoid extreme accelerations due to artificial stress. We are able to deal with the tensile instability without the need for XSPH.

With tSPH the cylinders shatter with all artificial stress treatments, as shown in Fig. 4.3. This is consistent with the need for XSPH to deal with the tensile instability with past, traditional SPH tests of colliding rubber cylinders. In all tSPH cases the cylinders shatter at earlier times than in the REMIX simulation with no artificial stress.

Next, we consider the relative strength of two different artificial stress schemes in the REMIX simulations. For each interaction between a pair of particles, we consider the fractional contribution of the artificial stress relative to the unmodified stress tensor,

$$\Delta_{ij} \equiv \frac{\parallel \sigma_{\text{art},\,ij} \parallel}{\parallel \sigma_i \parallel} \;, \tag{4.34}$$

where the norm of the tensors is given by $\parallel \sigma_i \parallel \equiv \sqrt{\sigma_i^{\alpha\beta} \sigma_i^{\alpha\beta}}$ and similarly for $\sigma_{\text{art},\,ij}$. For the two methods, $\sigma_{\text{art},\,ij}$ is calculated by either Eqn. 4.12 or 4.15. To get a measure of the strength of the correction applied to a particle, we calculate the average value of Δ_{ij} for each interaction with neighbouring particles,

$$\bar{\Delta}_i \equiv \sum_j^{N_j} \frac{\Delta_{ij}}{N_j} \;, \tag{4.35}$$

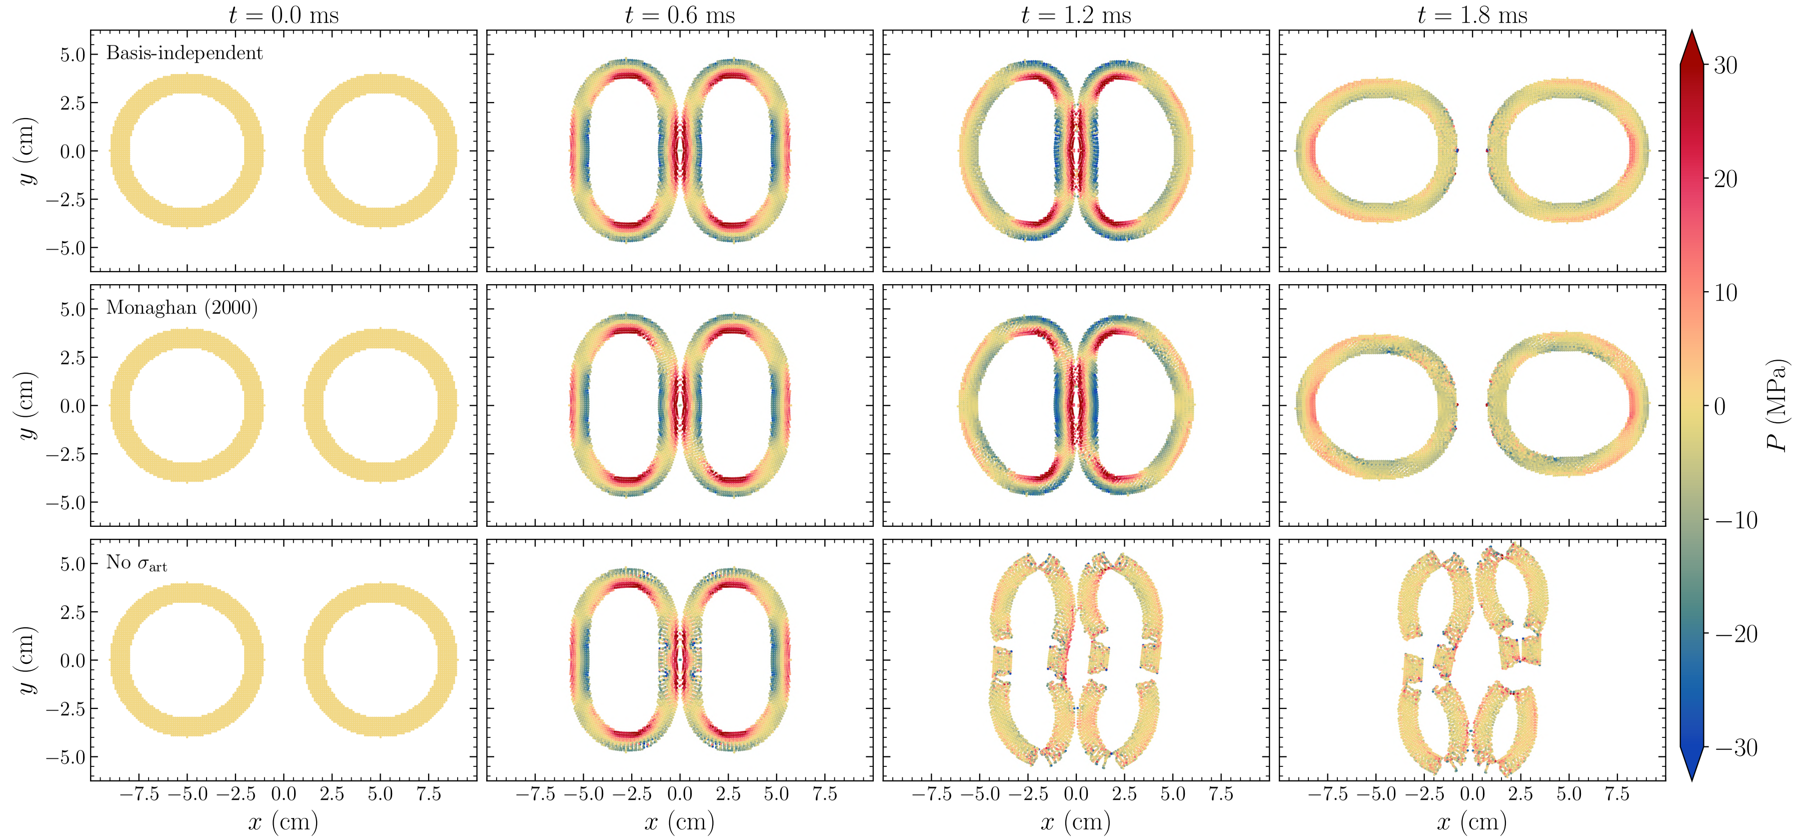

Figure 4.2: Snapshots from REMIX simulations of colliding rubber cylinders. Columns show four times and rows correspond to three different simulations carried out with: our new basis-independent artificial stress; the Monaghan (2000) artificial stress; and no artificial stress. Individual particles are plotted and coloured by their pressure.

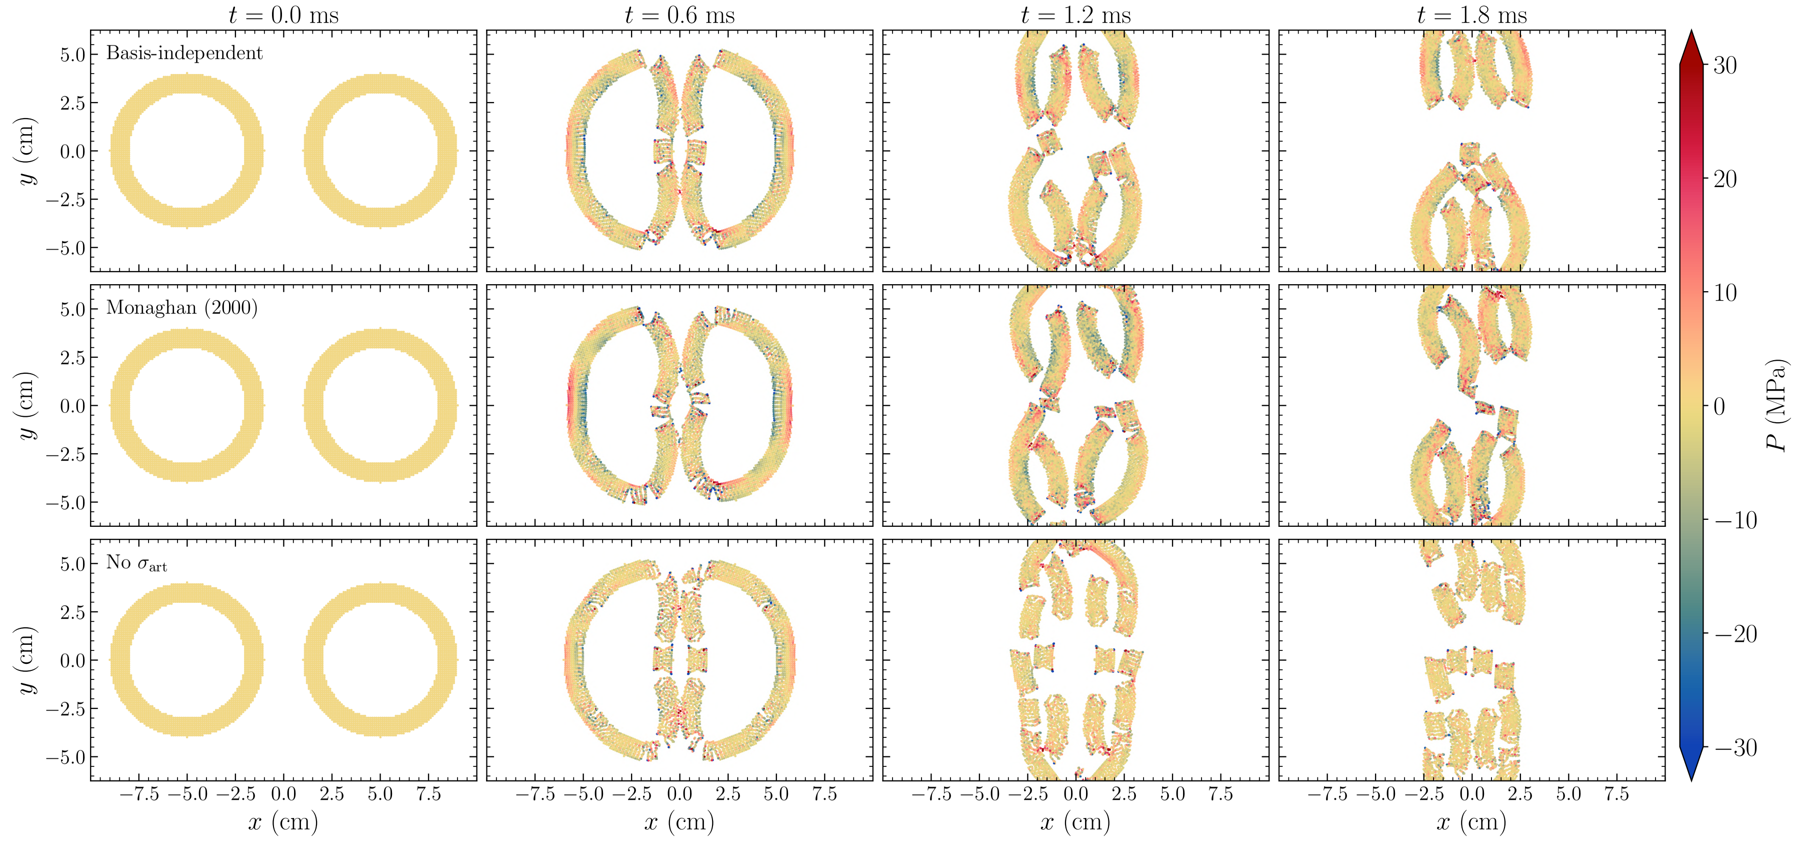

Figure 4.3: Snapshots from traditional SPH (tSPH) simulations of colliding rubber cylinders. Columns show four times and rows correspond to three different simulations carried out with: our new basis-independent artificial stress; the Monaghan (2000) artificial stress; and no artificial stress. Individual particles are plotted and coloured by their pressure.

where N_j is the number neighbours j for particle i.

Histograms of the $\bar{\Delta}_i$ of all simulation particles are shown in Fig. 4.4(a–c) at three times in the REMIX simulations that either use the basis-independent artificial stress approach or the M00 approach. We see that with the M00 approach there is consistently a large spike of particles with $\bar{\Delta}_i \sim 0.01$ (the relative size of which is more clearly shown in the inset with linear axes), whereas with our approach the number of particles in each bin steadily decreases with increased $\bar{\Delta}$. Additionally, with the M00 approach there is a relatively small, yet not insignificant, number of particles that have substantially higher values of $\bar{\Delta}_i$, as high as 0.15. With our approach individual particles do not experience these strong corrections and 3.2–4.4 times as many particles experience no correction at all at the three times shown in the figure.

In Fig. 4.4(d–f), we plot the cumulative sum of $\bar{\Delta}_i$ to investigate the strength of the overall correction in the simulation and the contribution to this from individual corrections of different magnitudes. Across these times, this measure of the overall correction is consistently ∼3 times higher with the M00 artificial stress than with our basis independent approach over these three times. The majority of the overall M00 correction is due to individual particles in the $\bar{\Delta}_i \sim 0.01$ spike. With our approach, the cumulative measure of $\bar{\Delta}_i$ increases close to linearly between $\bar{\Delta}_i = 0$ and 0.02, at which point it flattens out. Therefore, the overall strength of the artificial stress correction is not only not as large as with the M00 artificial stress, but the individual contributions that build it up are less strong than those that dominate the M00 correction.

Despite the correction being less strong than the M00 artifical stress, our approach is able to effectively address the tensile instability that would cause the shattering of the cylinders due to spurious fractures. This approach is not only more targeted but, more importantly, is not sensitive to the coordinate system used to calculate it.

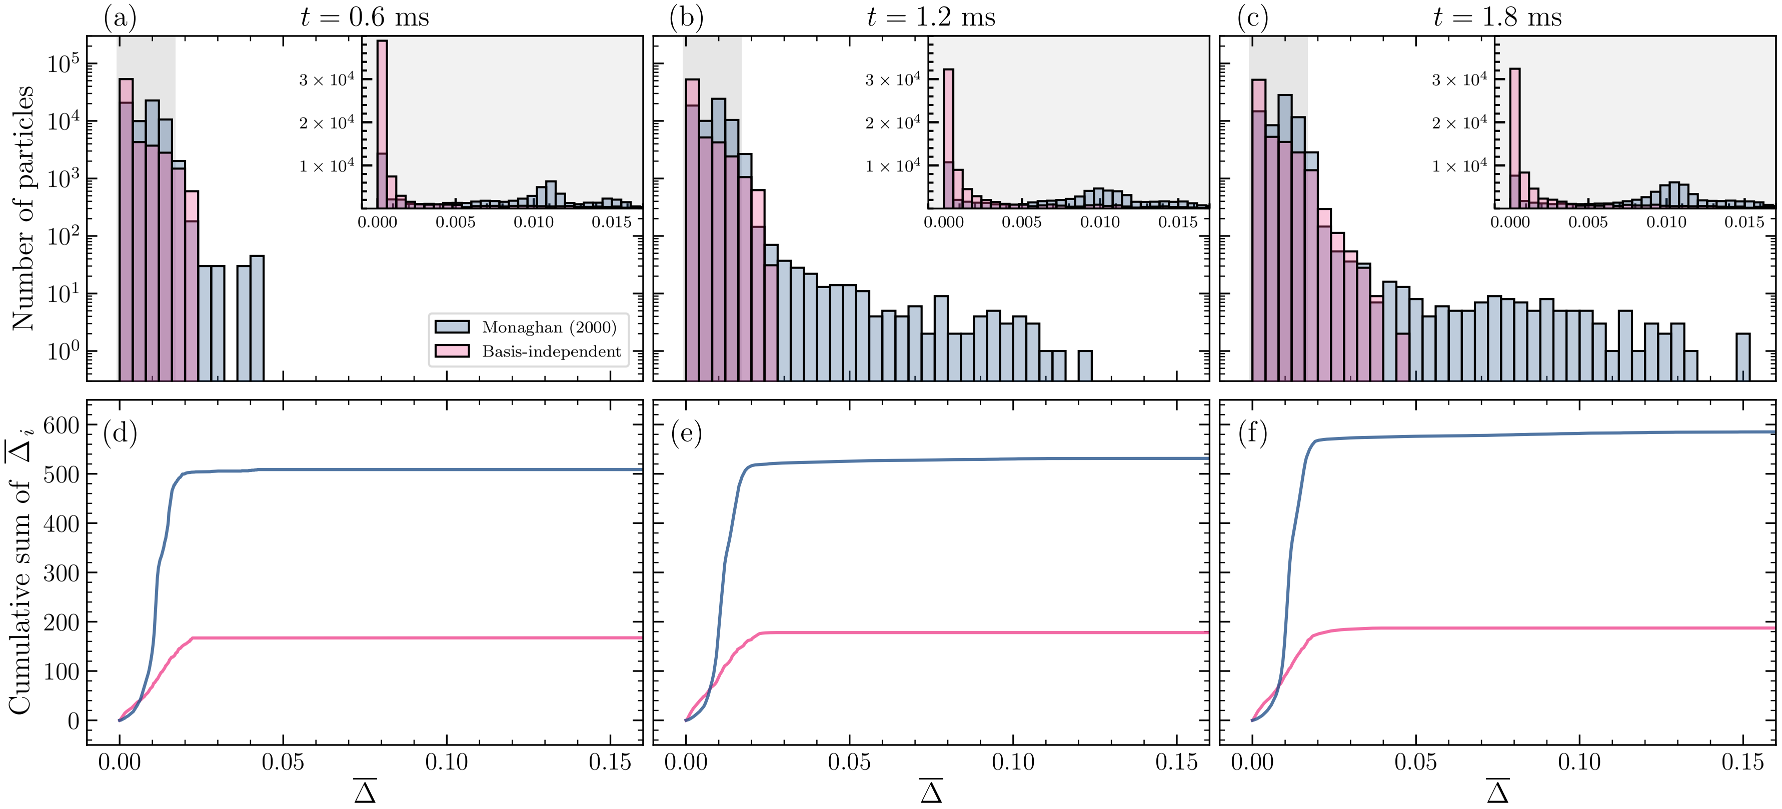

Figure 4.4: Plots showing the comparative strength of our basis-independent artificial stress approach and the Monaghan (2000) artificial stress at three times in REMIX simulations of colliding cylinders. Panels (a–c) show histograms of the number of particles with different $\bar{\Delta}_i$. Insets show equivalent histograms with a linear y-axis and for a smaller range of $\bar{\Delta}_i$, corresponding to the shaded region. Panels (d–f) show the cumulative sum of $\bar{\Delta}_i$ with the two methods – the sum of the $\bar{\Delta}_i$ of all particles with $\bar{\Delta}_i$ less than the x-axis value.

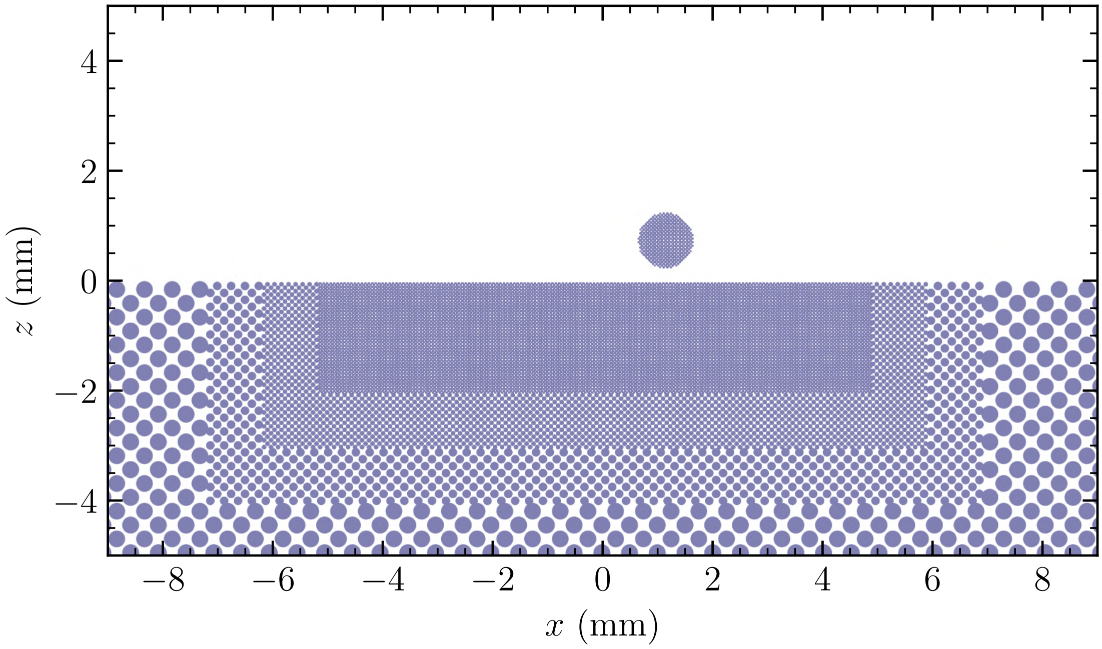

Figure 4.5: Cross-section from the initial conditions of an impact into aluminium at 50°, illustrating the variable resolution surface. Plotted particle radii are scaled by the cube root of their mass.

4.3.2 Impact cratering into aluminium

In this section, we use impacts into aluminum to test REMIX with the models for plastic deformation, described in §4.2.5. We compare the post-impact crater shapes against those formed in the equivalent experiments of Burchell & Mackay (1998) and the iSALE simulations of Davison et al. (2011). As well as the models for plasticity, these impacts use the methods for: fluid–solid interactions (§4.2.2), elasticity (§4.2.3), the basis-independent artificial stress formulation (§4.2.4), thermal weakening (§4.2.5), and density weakening (§4.2.5). Here we use a yield criterion given by Eqn. 4.23, despite using a constant yield stress, to match the methods used by Davison et al. (2011).

Both the target surface and impactor use a Tillotson equation of state for aluminium, the parameters of which are listed in Table 4.1, along with the material strength parameters used for the simulations. The density of both the impactor and the surface is initialised with $\rho = 2700$ kg m^{-3}. Particles are initialised in body-centred cubic (BCC) lattices such that the impactor and a region of dimensions 10 mm × 10 mm × 2 mm at the impact surface are high-resolution (where

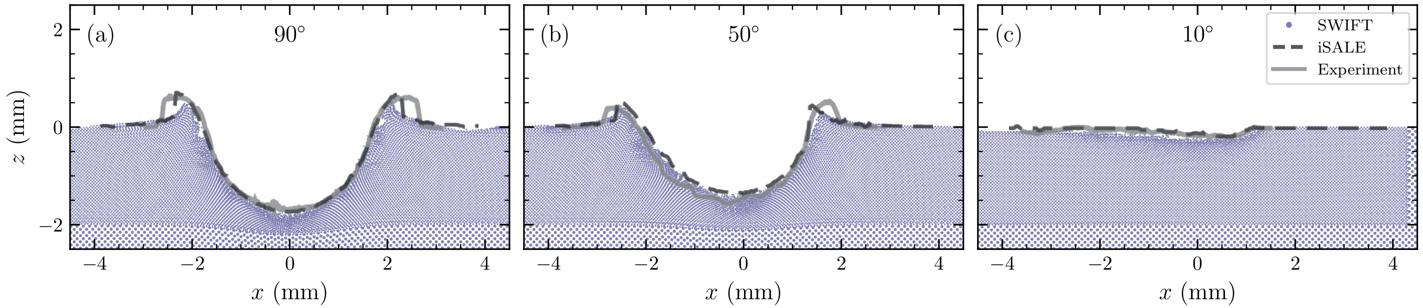

Figure 4.6: Cross-sections from REMIX simulations of impacts into aluminium, at $t = 50\ \mu$s after impact. Panels correspond to different simulations with three impact angles. The dashed line shows the crater formed in iSALE simulations and the solid line shows the experimental crater shape. Only particles with $\rho > 0.85\rho_0$ are plotted to isolate the crater from the remaining low-density material.

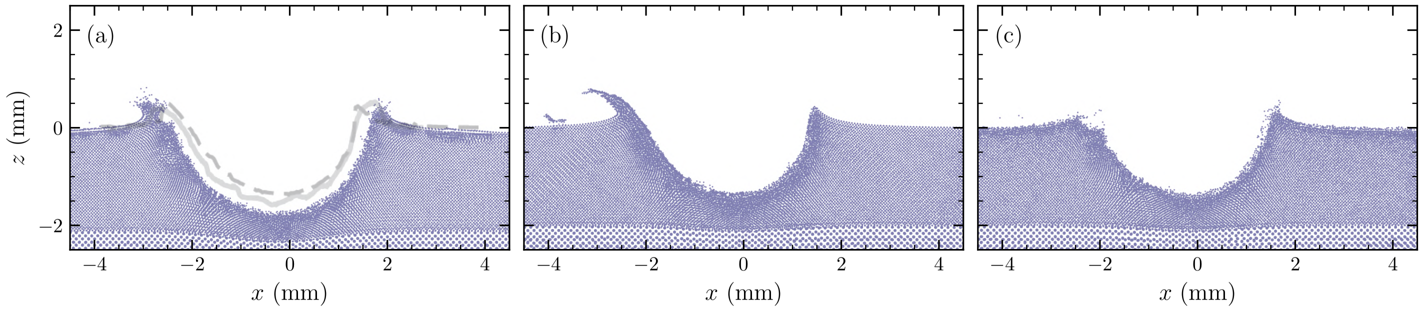

Figure 4.7: Cross-sections from REMIX simulations of impacts into aluminium, at $t = 50\ \mu$s after impact. Panels correspond to different simulations with three variations in methods: (a) with a yield criterion given by Eqn. 4.19 rather than Eqn. 4.23; (b) without limiting negative pressures by $-Y_0$; (c) without density weakening (Eqn. 4.25). All impacts are at an angle 50°. In (a), the dashed line and solid lines shows the crater formed in iSALE simulations and experiments, to illustrate the increased depth of the crater with this yield criterion.

the z-direction is perpendicular to the surface), sampled by many low-mass particles. The high resolution region of the surface is surrounded by a transition region in which the resolution decreases in discrete steps of 1/8 times as many particles per volume, until the low-resolution region with 1/512 times as many particles per volume. The full surface has dimensions 40 mm × 40 mm × 10 mm and consists of a large region of low-resolution material so that the shock produced by the impact does not affect the shape of the crater after reflection at the edge of the simulation domain. The transition region is required to prevent unstable configurations with larger steps in particle mass. For the impacts presented here, we use a resolution of 10 particles-per-projectile-radius, calculated based on the equivalent cubic lattice with the same number of particles per volume as the BCC lattice we use. A central cross-section of the initial impact configuration for an impact at 50° to the surface is shown in Fig. 4.5.

The resulting craters from impacts at three angles are shown in Fig. 4.6. The shapes of the craters produced using REMIX closely follow those produced in the iSALE simulations and the experiments. They follow the iSALE craters more closely than the experiments, which is perhaps not surprising, since both simulation methods use the same physical models for plastic deformation.

Fig. 4.7 shows the shapes of craters formed from 50° impacts using slight variations in numerical methods. In (a) we show the crater formed when a yield criterion given by Eqn. 4.19 is used instead of Eqn. 4.23. This leads to plastic deformation at smaller yields, and so the material is more soft and the crater formed is slightly deeper. In (b) we show the results of not limiting the negative pressures by the yield stress at zero pressure. The large negative pressures in this scenario lead to the accumulation of material at the rim of the crater. In contrast, if no density weakening was used, as is the case in (c), then the material at the crater rim gets removed by the low-density material moving over it during the impact.

4.3.3 Impact into a basalt sphere

In this section we use an impact into a brittle, basalt sphere to test the fracture model described in §4.2.6. The initial conditions are based on the experiments carried out

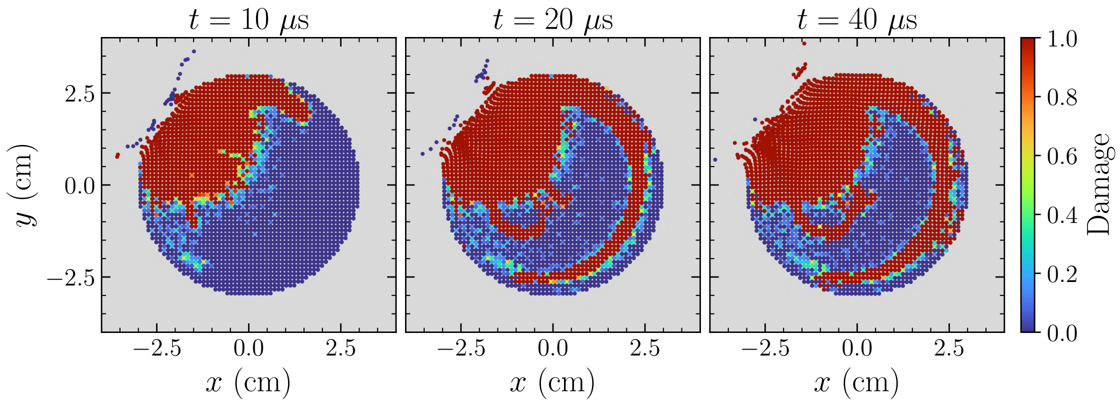

Figure 4.8: Cross-sections showing the accumulation of damage in REMIX simulations of an impact into a basalt sphere with $m_W = 8.5$. Panels show three times.

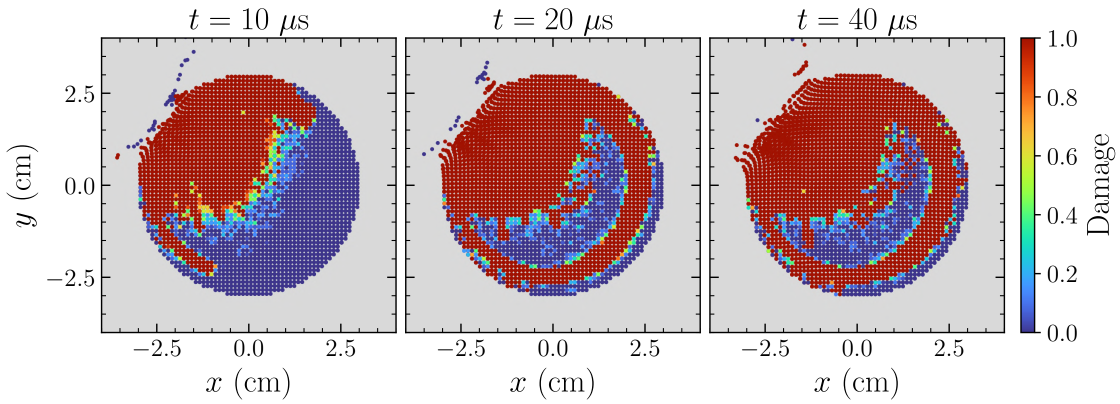

Figure 4.9: Cross-sections showing the accumulation of damage in REMIX simulations of an impact into a basalt sphere with $m_W = 8$. Panels show three times.

by Nakamura & Fujiwara (1991) and the simulations carried out by Benz & Asphaug (1994, 1995); Schäfer et al. (2016). For this test, we use the methods described in §4.2.3, §4.2.4, §4.2.5, and §4.2.6. We carry out separate tests with a constant yield stress (§4.2.5) or with a pressure-dependent yield stress that can be affected by damage (§4.2.5).

The particles in the basalt sphere use the Tillotson basalt EoS and the impactor is treated as lucite (Schäfer et al., 2016), the parameters for both of which are given in Table 4.1. Basalt particles are placed in a cubic lattice in a sphere of radius 3 cm with masses chosen to give a density 2700 kg m^{-3}. Here we use a resolution of 92,576 SPH particles, to closely follow the initial conditions of Benz & Asphaug (1995). The impactor uses particles of the same mass as the basalt sphere and has a total mass 0.2 g and density 1180 kg m^{-3}. The impactor has speed 3.2 km s^{-1}

and impact angle 30°. Particle internal energies are chosen to give low pressures corresponding to the 400 Pa with which the experiments of Nakamura & Fujiwara (1991) were carried out.

We present results from simulations with both the Weibull distribution parameters of Benz & Asphaug (1994) ($k_W = 5 \times 10^{28}$ cm^{-3} and $m_W = 8.5$) and those of Schäfer et al. (2016) ($k_W = 5 \times 10^{28}$ cm^{-3} and $m_W = 8$). The choice of m_W here leads to flaws in the material being seeded with different activation energies, given by Eqn. 4.26. The total number of flaws is expected to be $N_\mathrm{p} \ln N_\mathrm{p} = 10^6$, where N_p is the total number of particles (Weibull, 1939). Since $10^6/(k_W V) \sim 10^{-25} < 1$, a higher value of m_W increases the failure strains of all flaws, when keeping k_W constant.

Flaws are seeded in the basalt sphere only. The results of the impacts with these two sets of Weibull parameters are are shown in Fig. 4.8, for $m_W = 8.5$, and Fig. 4.9, for $m_W = 8$. The area surrounding the impact location gets damaged quickly after impact. Regions at the surface of the sphere get damaged by spallation stresses as compressive waves reflect at the free surface (Schäfer et al., 2016) and protruding fractures grow at the edges of the fully-damaged regions. With a lower value of m_W, lower strains are required to activate fracturing in each flaw, so in this case the basalt sphere is more damaged.

The overall growth of fractures shows strong similarities with the simulations of Benz & Asphaug (1994, 1995) and Schäfer et al. (2016). Future work involves the use of a friends-of-friends algorithm to analyse the size distributions of fragments produced by these impacts, and to compare them with the experimental results of Nakamura & Fujiwara (1991). Additionally, this will enable direct quantitative comparisons with the simulations of Benz & Asphaug (1994, 1995) and Schäfer et al. (2016) to help understand how the effect of the choice of m_W on fracture growth compares between different simulation methods.

4.4 Conclusions

The implementation of models for material strength into REMIX extends its applications to an entirely different regime of planetary impact simulations. We have presented a range of physical models and numerical treatments that combine to allow for the simulation of elasticity, plastic deformation and the growth of fractures.

In simulations of collisions between perfectly elastic rubber cylinders, we show that, with artificial stress, REMIX is able to simulate an appropriate collision, without needing to fundamentally modify the core approach of our methods, as is often done by using XSPH. We present a new, basis-independent approach to artificial stress that addresses the spurious SPH tensile instability in a way that is not only independent of the coordinate system in which the calculations are performed, in contrast with some past approaches, but that is more targeted in its correction. We show that this more targeted approach leads to the stress tensor of fewer particles being corrected and the corrections that are applied are typically smaller in magnitude.

I present results from simulations of crater-forming impacts onto ductile aluminium surfaces. The shape of the crater produced in REMIX simulations closely follows the shapes produced in equivalent experiments and iSALE simulations. However, we find that the crater shape, and in particular the shape of the crater rim, are sensitive to specifics of the numerical methods used to represent the physical processes at play. For future work, we will continue to investigate the SPH treatment of fluid–solid interactions and plastic deformation under tension to further develop our methods to treat these and similar scenarios in a way that is most physically appropriate.

Simulations of impacts into a brittle, basalt sphere are used to test the implementation of the model for tensile fracturing. We tested the effect of the parameters used to seed flaws in the material, verifying that lower activation thresholds lead to more fracturing of the brittle material. Future work involves the use of a friends-of-friends algorithm to determine the size distribution of fragments produced in this experiment and comparing these with experimental results and past SPH simulations.

The validation of material strength physics models is an important step for the development of the REMIX SPH formulation. With the improved numerical treatment of REMIX, compared with traditional SPH formulations, and the state-of-the-art computational efficiency of the Swift code, once fully validated, these methods will allow simulations of planetary impacts across a range of scales, where the effects of solid-body mechanics become significant, at high resolutions and to high accuracy.

CHAPTER 5

Conclusions

Planetary impacts shape the evolution of planetary systems, from the craters left by asteroids and comets on the surfaces of planetary bodies to the catastrophic giant impacts that dominate the final stages of planet formation. The footprints of these impacts still persist throughout the Solar System and so performing numerical simulations of these impact scenarios allows us to understand the processes that formed our Solar System, as well as planetary systems around distant stars.

In this thesis I have developed and applied novel numerical methods that significantly improve the treatment of physical processes that are key to simulations of planetary impacts. Through extensive validation of our methods we are able to tackle unanswered questions with a higher confidence in the accuracy of our simulations. I have tested these methods in simulations of giant impacts onto Jupiter to investigate whether a single giant impact could be responsible for the planet's dilute core. Additionally, I have extended our newly developed simulation methods to include a range of models for solid-body material strength physics.

In Chapter 2 I presented REMIX, a smoothed particle hydrodynamics (SPH) scheme designed to alleviate effects that typically suppress mixing and instability growth at density discontinuities in SPH simulations. This problem was approached by directly targeting sources of kernel smoothing error and discretisation error, resulting in a generalised, material-independent formulation that improves the treatment both of discontinuities within a single material and of interfaces between dissimilar materials. REMIX is built upon a fully compressible and thermodynamically consistent core-SPH construction, retaining Galilean invariance as well as conservation of mass, momentum, and energy, and is integrated in the open-source, state-of-the-art SWIFT code. Designed with computational efficiency in mind, REMIX can be used for high-resolution simulations without significant cost to run-speed, as demonstrated throughout this thesis.

I have validated REMIX in a selection of test simulations, chosen to investigate the performance of REMIX in capturing a range of hydrodynamic behaviour that is important in simulations of planetary impacts. REMIX demonstrates marked improvements compared with a traditional SPH formulation, in particular in tests focusing on the treatment of mixing. For the first time, REMIX allows appropriate treatment of material mixing in SPH simulations, without introducing additional equation of state dependence in, for example, particle volume elements; without requiring bespoke choices of particle masses; and without contrived or targeted corrections.

Next, in Chapter 3, I applied REMIX to simulations of giant impacts onto Jupiter to investigate whether the planet's dilute core could be formed from a single giant impact. After demonstrating the improvements in the treatment of mixing in this regime with REMIX compared with a traditional SPH formulation, I carried out simulations at a range of impact speeds and angles, at different numerical resolutions, and with different internal structures and equations of state. None of our simulations produce a dilute core immediately after impact, in contrast with what had been suggested was possible by past work. Therefore, our simulations suggest that it is unlikely that a single giant impact could be responsible for Jupiter's dilute core, and that the result of the previous study is an artefact of the numerical methods that were employed.

In Chapter 4, I presented the physical models implemented into the REMIX scheme to elevate it from a purely hydrodynamic method to one that can capture material strength properties. This includes treatments of elasticity, validated in simulations of collisions between rubber cylinders; plasticity, validated in crater-forming impacts into aluminium with comparisons against experimental results and contemporary numerical methods; and the growth of fractures, tested in simulations of impacts onto a basalt sphere. I presented a new, basis-independent artificial stress treatment and found that it is effective in targeting problematic particles in simulations of colliding rubber cylinders. In these simulations REMIX is able to accurately simulate the flexing of the cylinders without forming spurious fractures caused by the tensile instability. This is not the case for traditional SPH formulations that fail this test unless additional corrections are applied.

Future work involves the further validation of the newly developed methods presented here, including the treatment of demixing in REMIX simulations, such as in the giant impacts onto Jupiter presented here, to investigate whether we physically capture these processes or whether an extension of our methods that includes mixing in the sub-particle scale would be required. Additionally, it would be beneficial to further investigate the treatment of vacuum interfaces in typical giant impact simulations, where accurately simulating the thermodynamics of material that is ejected during the impact is often important. The continued validation of REMIX with material strength physics would require testing in additional scenarios, such as to investigate the treatment of fluid–solid interactions; in crater-forming impacts into targets more representative of planetary surfaces; and in giant impacts.

REMIX opens up a wide range of possible science applications, simulations for which can now be performed with significantly improved accuracy. This includes investigating mixing processes during the giant impact that formed the Moon, an unsolved problem that continues to garner the interest of planetary scientists. Additionally, with material strength physics implemented with REMIX we can now investigate a new regime of impacts, from small crater-forming impacts onto the Moon, to larger impacts onto Mars.

Planetary science is entering an exciting new age, as rapid advancements are allowing observations of more and more planetary systems around distant stars. Along with these observations come new theories to be tested and a need for sophisticated numerical methods that can simulate the processes that govern the evolution of planetary systems. Here I have presented newly developed methods to perform hydrodynamic simulations that show significant advancements in the treatment of both fluid and solid-body dynamics, in particular related to material mixing. These improvements open up a range of science applications that can help to further our understanding of the dynamics and details of planetary impacts and the evolution of planetary systems.

Bibliography

ALMA Partnership et al., 2015, , 808, L3 1

Agertz O., et al., 2007, Monthly Notices of the Royal Astronomical Society, 380, 963 1.2, 2.3.3, 2.3.3, 2.3.7, 2.3.7, 3.1

Antuono M., Colagrossi A., Marrone S., Molteni D., 2010, Computer Physics Communications, 181, 532 2.2.1, 2.2.6

Antuono M., Colagrossi A., Marrone S., 2012, Computer Physics Communications, 183, 2570 2.2.6

Armitage P. J., 2020, Astrophysics of planet formation. Cambridge University Press 1

Asphaug E., Benz W., 1996, Icarus, 121, 225 1.2, 4.1

Baldwin R. B., 1949, Chicago] Univ. of Chicago Press [1949 1

Ballantyne H. A., Jutzi M., Golabek G. J., Mishra L., Cheng K. W., Rozel A. B., Tackley P. J., 2023, Icarus, 392, 115395 1.2, 4.1

Balsara D. S., 1995, Journal of Computational Physics, 121, 357 2.2.5

Bauer A., Springel V., 2012, Monthly Notices of the Royal Astronomical Society, 423, 2558 2.3.7

Becker A., Lorenzen W., Fortney J. J., Nettelmann N., Schöttler M., Redmer R., 2014, The Astrophysical Journal Supplement Series, 215, 21 3.2.2

Benz W., 1988, Computer Physics Communications, 48, 97 1.2

Benz W., 2000, in From Dust to Terrestrial Planets: Proceedings of an ISSI Workshop, 15–19 February 1999, Bern, Switzerland. pp 279–294 1

Benz W., Asphaug E., 1994, Icarus, 107, 98 1.2, 4.1, 4.2.1, 4.2.5, 4.2.6, 4.3.3

Benz W., Asphaug E., 1995, Computer physics communications, 87, 253 1.2, 2.2.1, 4.1, 4.2.6, 4.2.6, 4.3.3

Benz W., Slattery W., Cameron A., 1986, Icarus, 66, 515 1

Benz W., Cameron A., Melosh H., 1989, Icarus, 81, 113 1.3.1

Borrow J., Schaller M., Bower R. G., Schaye J., 2022, Monthly Notices of the Royal

Astronomical Society, 511, 2367 2.2.5, 2.3.8, 2.15, C.3, C.3

Braspenning J., Schaye J., Borrow J., Schaller M., 2023, Monthly Notices of the Royal Astronomical Society, 523, 1280 2.3.7, 2.3.7, 2.3.7, 2.3.7

Burchell M. J., Mackay N. G., 1998, Journal of Geophysical Research: Planets, 103, 22761 4.3.2

Cameron A. G., Ward W. R., 1976, in Abstracts of the Lunar and Planetary Science Conference, volume 7, page 120,(1976). 1

Canup R. M., 2004, Icarus, 168, 433 2.2.5

Caupin F., Herbert E., 2006, Comptes Rendus Physique, 7, 1000 4.2.2

Chabrier G., Debras F., 2021, The Astrophysical Journal, 917, 4 2.1.1, 2.3.9, 3.2.2

Chabrier G., Mazevet S., Soubiran F., 2019, The Astrophysical Journal, 872, 51 3.1, 3.2.2

Chandrasekhar S., 1961, Hydrodynamic and hydromagnetic stability 2.3.3, 2.3.3, 2.3.5, 3.3.1, 3.3.2

Collins G. S., Melosh H. J., Ivanov B. A., 2004, Meteoritics & Planetary Science, 39, 217 4.1, 1, 4.2.5, 4.2.5, 4.2.6

Couch S. M., Graziani C., Flocke N., 2013, The Astrophysical Journal, 778, 181 3.1

Cullen L., Dehnen W., 2010, Monthly Notices of the Royal Astronomical Society, 408, 669 2.2.5

Davison T. M., Collins G. S., Elbeshausen D., Wuennemann K., Kearsley A., 2011, Meteoritics & Planetary Science, 46, 1510 4.3.2

Debras F., Chabrier G., 2019, The Astrophysical Journal, 872, 100 3.1

Dehnen W., Aly H., 2012, Monthly Notices of the Royal Astronomical Society, 425, 1068 1.2.1, 1.2.1, 1.3.2, 3.2.1, D

Deng H., Reinhardt C., Benitez F., Mayer L., Stadel J., Barr A. C., 2019, The Astrophysical Journal, 870, 127 1.3.1

Dienes J. K., 1979, Acta mechanica, 32, 217 4.2.3

Emsenhuber A., Jutzi M., Benz W., 2018, Icarus, 301, 247 1.2, 4.1, 4.2.5

Evrard A. E., 1988, Monthly Notices of the Royal Astronomical Society (ISSN 0035-8711), vol. 235, Dec. 15, 1988, p. 911-934., 235, 911 2.3.8

Fermi E., Pasta P., Ulam S., Tsingou M., 1955, Technical report, Studies of the nonlinear problems. Los Alamos National Laboratory (LANL), Los Alamos, NM (United States) 1

Frontiere N., Raskin C. D., Owen J. M., 2017, Journal of Computational Physics, 332, 160 1.3.2, 2.2.3, 2.2.3, 2.2.3, 2.2.3, 2.2.5, 2.2.5, 2.3.5, 2.3.7, 3.2.1, 3.3.1, F

Fryxell B., et al., 2000, The Astrophysical Journal Supplement Series, 131, 273 3.1

Gaburov E., Nitadori K., 2011, Monthly Notices of the Royal Astronomical Society, 414, 129 1.3.1

García-Senz D., Cabezón R. M., Escartín J. A., 2012, Astronomy & astrophysics, 538, A9 1.3.2

Gilbert G. K., 1893, The moon's face. A study of the origin of its features. 1

Gingold R. A., Monaghan J. J., 1977, Monthly notices of the royal astronomical society, 181, 375 1.2, 3.1

Goldreich P., Ward W. R., 1973, Astrophysical Journal, Vol. 183, pp. 1051-1062 (1973), 183, 1051 1

Grady D. E., Kipp M. E., 1980, in International journal of rock mechanics and mining sciences & geomechanics abstracts. pp 147–157 4.2.6

Gray J. P., Monaghan J. J., Swift R., 2001, Computer methods in applied mechanics and engineering, 190, 6641 4.2.1, 4.2.4, 4.2.4

Haisch Jr K. E., Lada E. A., Lada C. J., 2001, The Astrophysical Journal, 553, L153 1

Haldemann J., Alibert Y., Mordasini C., Benz W., 2020, Astronomy & Astrophysics, 643, A105 2.1.1, 2.3.9, 3.2.2

Hartmann W. K., Davis D. R., 1975, Icarus, 24, 504 1

Helled R., et al., 2022, Icarus, p. 114937 3.1

Holsapple K. A., Michel P., 2008, Icarus, 193, 283 4

Hopkins P. F., 2013, Monthly Notices of the Royal Astronomical Society, 428, 2840 1, 1.2.2, 2.2.3, 2.2.3, 2.3.3

Hopkins P. F., 2015, Monthly Notices of the Royal Astronomical Society, 450, 53 1.3.1, 2.3.3

Hosono N., Saitoh T. R., Makino J., 2013, Publications of the Astronomical Society of Japan, 65, 108 1.3.1

Hosono N., Saitoh T. R., Makino J., Genda H., Ida S., 2016, Icarus, 271, 131 1.3.1

Howard S., et al., 2023, Astronomy & Astrophysics, 672, A33 3.1, 3.2.2

Inutsuka S.-i., 2002, Journal of Computational Physics, 179, 238 1.3.1

Ivanova N., et al., 2013, The Astronomy and Astrophysics Review, 21, 1 1.3.1

Jutzi M., 2015, Planetary and space science, 107, 3 1.2, 4.1

Jutzi M., Benz W., Michel P., 2008, Icarus, 198, 242 1.2, 4.1

Kegerreis J., Eke V., Gonnet P., Korycansky D., Massey R., Schaller M., Teodoro L., 2019, Monthly Notices of the Royal Astronomical Society, 487, 5029 2.1.2, 2.3.9, 3.2.1, 3.4.3

Kegerreis J. A., Eke V. R., Catling D. C., Massey R. J., Teodoro L. F., Zahnle K. J., 2020, The Astrophysical Journal Letters, 901, L31 2.3.9

Kegerreis J. A., Ruiz-Bonilla S., Eke V. R., Massey R. J., Sandnes T. D., Teodoro L. F., 2022, The Astrophysical Journal Letters, 937, L40 2.3.9, 3.1

Kenworthy M., et al., 2023, , 622, 251 1

LeVeque R. J., 2007, Finite difference methods for ordinary and partial differential equations: steady-state and time-dependent problems. SIAM 1.1

Li S. L., Agnor C., Lin D., 2010, The Astrophysical Journal, 720, 1161 3.1

Libersky L. D., Petschek A. G., 1991, in Trease H. E., Fritts M. F., Crowley W. P., eds, , Vol. 395, Advances in the Free-Lagrange Method Including Contributions on Adaptive Gridding and the Smooth Particle Hydrodynamics Method. pp 248–257, doi:10.1007/3-540-54960-9$_5$8 1.2, 4.1, 4.2.3

Libersky L. D., Petschek A. G., Carney T. C., Hipp J. R., Allahdadi F. A., 1993, Journal of Computational Physics, 109, 67 1.2

Lind S. J., Rogers B. D., Stansby P. K., 2020, Proceedings of the Royal Society A, 476, 20190801 1.2

Liu W. K., Jun S., 1998, International Journal for Numerical Methods in Engineering, 41, 1339 1.3.2

Liu W. K., Jun S., Zhang Y. F., 1995, International Journal for Numerical Methods in Fluids, 20, 1081 1.3.2

Liu S.-F., Hori Y., Lin D., Asphaug E., 2015, The Astrophysical Journal, 812, 164 3.1

Liu S.-F., Hori Y., Müller S., Zheng X., Helled R., Lin D., Isella A., 2019, Nature, 572, 355 1.4, 3.1, 3.2.2, 3.2.3, 3.4, 3.4.4, 3.10, 3.11, 3.5, 3.6

Lucy L. B., 1977, Astronomical Journal, vol. 82, Dec. 1977, p. 1013-1024., 82, 1013 1.2, 3.1

Lundborg N., 1968, in International Journal of Rock Mechanics and Mining Sciences & Geomechanics Abstracts. pp 427–454 4.2.5

Luther R., et al., 2022, The Planetary Science Journal, 3, 227 4.2.5, 4.2.5, 4.2.5

Lyon S. P., 1978, LANL report 3.4.4

Mankovich C. R., Fuller J., 2021, Nature Astronomy, 5, 1103 3.1, 3.5

Mayor M., Queloz D., 1995, nature, 378, 355 1

Mazevet S., Licari A., Soubiran F., 2022, Astronomy & Astrophysics, 664, A112 3.2.2

McNally C. P., Lyra W., Passy J.-C., 2012, The Astrophysical Journal Supplement Series, 201, 18 2.3.3, 2.3.3, 2.3.3, 2.4, 2.3.3, 3.3.1

Melosh H. J., 1989, New York: Oxford University Press; Oxford: Clarendon Press 1.3.1, 2.1.1, 3.2.2, 4.2.2

Miguel Y., Guillot T., Fayon L., 2016, Astronomy & Astrophysics, 596, A114 3.2.2

Miguel Y., et al., 2022, Astronomy & Astrophysics, 662, A18 3.1

Militzer B., Hubbard W. B., 2013, The Astrophysical Journal, 774, 148 3.2.2

Militzer B., et al., 2022, The planetary science journal, 3, 185 3.1

Mimeau C., Mortazavi I., 2021, Fluids, 6, 68 1.1

Moll R., Garaud P., Mankovich C., Fortney J., 2017, The Astrophysical Journal, 849, 24 3.1

Monaghan J., 1985, Journal of Computational Physics, 60, 253 1.2, 1.3.1

Monaghan J., 1989, Journal of Computational physics, 82, 1 4.3.1

Monaghan J. J., 1992, In: Annual Review of Astronomy and Astrophysics. Vol. 30

(A93-25826 09-90), p. 543-574., 30, 543 1.2

Monaghan J. J., 1997, Journal of Computational Physics, 136, 298 2.2.6

Monaghan J. J., 2000, Journal of computational physics, 159, 290 4.2.4, 4.1, 4.2.4, 4.2.4, 4.3.1, 4.3.1, 4.2, 4.3, 4.3.1, 4.4

Monaghan J. J., Gingold R. A., 1983, Journal of Computational Physics, 52, 374 2.2.5, 2.2.5

Morris J. P., Monaghan J. J., 1997, Journal of Computational Physics, 136, 41 2.2.5

Moukalled F., Mangani L., Darwish M., Moukalled F., Mangani L., Darwish M., 2016, The finite volume method. Springer 1.1

Müller S., Helled R., Cumming A., 2020, Astronomy & Astrophysics, 638, A121 3.1

Nakamura A., Fujiwara A., 1991, Icarus, 92, 132 4.3.3

Nettelmann N., 2017, Astronomy & Astrophysics, 606, A139 3.1

Ohnaka M., 1995, Geophysical Research Letters, 22, 25 4.2.5

Pearl J. M., Raskin C. D., Owen J. M., 2022, Journal of Computational Physics, 469, 111533 1.3.1, 2.2.5, 2.2.6, 4.2.4

Price D. J., 2008, Journal of Computational Physics, 227, 10040 1.3.1, 2.2.6, 2.3.3, 2.3.3, 3.3.1, 3.3.1

Price D. J., 2012, Journal of Computational Physics, 231, 759 1.2, 1.2.1, 1.3.1, 1.3.1, 1.3.2, 2.2.5, 3.2.1, 3.2.1

Price D. J., et al., 2018, Publications of the Astronomical Society of Australia, 35, e031 2.2.3, 2.2.6, 2.2.6, F

Raducan S. D., Jutzi M., 2022, The planetary science journal, 3, 128 1.2, 4.1

Read J., Hayfield T., 2012, Monthly Notices of the Royal Astronomical Society, 422, 3037 2.2.5, 2.2.6

Read J., Hayfield T., Agertz O., 2010, Monthly Notices of the Royal Astronomical Society, 405, 1513 1.2, 1.2.2, 1.2.2, 1.3.1, 1.3.2, 2.2.2, 2.2.2, 2.3.3, 3.2.1

Reinhardt C., Stadel J., 2017, Monthly Notices of the Royal Astronomical Society, 467, 4252 1.3.1, 2.2.4, 2.2.5

Reinhardt C., Chau A., Stadel J., Helled R., 2020, Monthly Notices of the Royal Astronomical Society, 492, 5336 1.3.1

Ritchie B. W., Thomas P. A., 2001, Monthly Notices of the Royal Astronomical Society, 323, 743 1.3.1, 2.2.2

Robertson B. E., Kravtsov A. V., Gnedin N. Y., Abel T., Rudd D. H., 2010, Monthly Notices of the Royal Astronomical Society, 401, 2463 2.3.3, 3.1, 3.3.1, G

Rosswog S., 2015, Monthly Notices of the Royal Astronomical Society, 448, 3628 1.3.2

Rosswog S., 2020, Monthly Notices of the Royal Astronomical Society, 498, 4230 1.3.2, 2.2.5, 2.2.5, 2.2.6, 2.2.6, 2.3.1, 2.3.1, 2.3.3, 2.3.3, 3.3.1, F, F

Ruiz-Bonilla S., Eke V. R., Kegerreis J. A., Massey R. J., Teodoro L. F., 2021, Monthly Notices of the Royal Astronomical Society, 500, 2861 2.3.9, 3.1, 3.2.1

Ruiz-Bonilla S., Borrow J., Eke V. R., Kegerreis J. A., Massey R. J., Sandnes T. D., Teodoro L. F., 2022, Monthly Notices of the Royal Astronomical Society, 512, 4660 1.3.1

Safronov V. S., 1972, Evolution of the protoplanetary cloud and formation of the earth and planets. 1

Saitoh T. R., Makino J., 2013, The Astrophysical Journal, 768, 44 1.3.1, 2.3.1

Saumon D., Chabrier G., van Horn H. M., 1995, Astrophysical Journal Supplement v. 99, p. 713, 99, 713 3.2.2

Schäfer C., Riecker S., Maindl T. I., Speith R., Scherrer S., Kley W., 2016, Astronomy & Astrophysics, 590, A19 4.2.1, 4.2.4, 4.2.6, 4.1, 4.3.3

Schaller M., Dalla Vecchia C., Schaye J., Bower R. G., Theuns T., Crain R. A., Furlong M., McCarthy I. G., 2015, Monthly Notices of the Royal Astronomical Society, 454, 2277 2.2.6

Schaller M., Gonnet P., Chalk A. B. G., Draper P. W., 2016, in Proceedings of the Platform for Advanced Scientific Computing Conference. p. 2 (`arXiv:1606.02738`), doi:10.1145/2929908.2929916 2.1.2

Schaller M., et al., 2024, Monthly Notices of the Royal Astronomical Society, 530, 2378 2, 2.1.2, 3.1

Sigalotti L. D. G., Rendón O., Klapp J., Vargas C. A., Cruz F., 2019, Applied Mathematics and Computation, 356, 50 1.3.1

Sod G. A., 1978, Journal of Computational Physics, 27, 1 2.3.2

Spreng F., Vacondio R., Eberhard P., Williams J. R., 2020, Computers & Fluids, 198, 104388 1.3.2, 3.2.1

Springel V., 2005, Monthly Notices of the Royal Astronomical Society, 364, 1105 1.2

Springel V., 2010a, Annual Review of Astronomy and Astrophysics, 48, 391 1.2, 2.3.1

Springel V., 2010b, Monthly Notices of the Royal Astronomical Society, 401, 791 1.1, 2.3.8, 3.1

Stewart S., et al., 2020, in AIP conference proceedings. 2.1.1, 2.3.4, 2.3.9, 3.2.2

Sun P.-N., Le Touzé D., Oger G., Zhang A.-M., 2021, Journal of Computational Physics, 426, 109937 2.2.6

Tskhakaya D., Matyash K., Schneider R., Taccogna F., 2007, Contributions to Plasma Physics, 47, 563 1.1

Van Leer B., 1974, Journal of computational physics, 14, 361 2.2.5

Vazan A., Helled R., Guillot T., 2018, Astronomy & Astrophysics, 610, L14 3.1

Venturini J., Helled R., 2020, Astronomy & Astrophysics, 634, A31 3.1

Vila J. P., 1999, Mathematical Models and Methods in Applied Sciences, 9, 161 1.3.1

Violeau D., Fonty T., 2019, Computers & Fluids, 191, 104240 1.3.1, 3.2.1

Violeau D., Rogers B. D., 2016, Journal of Hydraulic Research, 54, 1 1.2

VonNeumann J., Richtmyer R. D., 1950, Journal of Applied Physics, 21, 232 2.2.5

Wadsley J. W., Keller B. W., Quinn T. R., 2017, Monthly Notices of the Royal Astronomical Society, 471, 2357 1.3.2, 2.2.2

Wahl S. M., et al., 2017, Geophysical Research Letters, 44, 4649 1.4, 3.1

Wang Z.-B., Chen R., Wang H., Liao Q., Zhu X., Li S.-Z., 2016, Applied Mathematical Modelling, 40, 9625 1.2

Weibull W., 1939, IVB-Handl. 4.2.6, 4.3.3

Wendland H., 1995, Advances in Computational Mathematics, 4, 389 1.2.1, 3.2.1, 3.2.4

Wilson H. F., Militzer B., 2011, The Astrophysical Journal, 745, 54 3.3, 3.5

Wilson H. F., Militzer B., 2012, Physical Review Letters, 108, 111101 3.3, 3.5

Woolfson M., 2007, Monthly Notices of the Royal Astronomical Society, 376, 1173 1.3.1

Youdin A. N., Shu F. H., 2002, The Astrophysical Journal, 580, 494 1

Zienkiewicz O. C., Taylor R. L., 2005, The finite element method set. Elsevier 1.1

APPENDIX A

Notation

$A,\ \mathbf{A}$	Scalar, vector
$A^{\alpha,\cdots,\omega}$	Elements of a vector, matrix, or higher-order tensor. Greek letter superscripts correspond to spatial dimensions (e.g. the x, y, and z components of a 3D vector) and like indices are summed over
$A_i,\ A_j$	Quantity associated with, or sampled at the position of: particle i; a neighbour j of particle i
A_{ij}	Pairwise interaction associated with neighbour j acting on particle i
$\mathbf{r}$	Position vector at which we probe the continuum fluid
$\mathbf{r}'$	Position vector integrated over in the continuum limit, for convolutions with a kernel function
$\mathbf{r}_i$	Position vector of particle i. When discretising the fluid, the notation changes $\mathbf{r} \Rightarrow \mathbf{r}_i$, as we probe the fluid at particle positions

$\mathbf{r}_{ij}$	$\mathbf{r}_i - \mathbf{r}_j$
m_i	Mass of particle i
V_i	Volume element of particle i
H	Kernel compact support. The radial extent from the kernel centre above which a kernel function is zero
h	Kernel smoothing length. We define the smoothing length as twice the standard deviation of the kernel
η_{kernel}	Parameter that scales the radial extent of the kernel
W_{ij}	Kernel function centred at $\mathbf{r}_i$, sampled at $\mathbf{r}_j$. $W_{ij} \equiv W(\mathbf{r}_{ij},\, h_i) \equiv W(\mathbf{r}_i - \mathbf{r}_j,\, h(\mathbf{r}_i))$
$\overline{W}_{ij}$	Symmetrised kernel: $\overline{W}_{ij} \equiv \overline{W}(\mathbf{r}_i - \mathbf{r}_j,\, h(\mathbf{r}_i),\, h(\mathbf{r}_j)) \equiv \frac{1}{2}\left(W(\mathbf{r}_{ij},\, h_i) + W(\mathbf{r}_{ji},\, h_j)\right)$
$\hat{W}_{ij}$	Kernel W_{ij} normalised to volume elements V_j
$\mathcal{W}_{ij}$	Linear-order reproducing kernel, constructed using $\overline{W}_{ij}$
$\widetilde{\mathcal{W}}_{ij}$	Linear-order reproducing kernel with vacuum boundary treatment, constructed using $\overline{W}_{ij}$
$\langle A_i \rangle$, $\overline{A}_i$, $\hat{A}_i$	Quantity calculated by kernel interpolation using W_{ij}, $\overline{W}_{ij}$, $\hat{W}_{ij}$, respectively
$\frac{dA}{d\mathbf{r}}$, $\frac{dA}{dr^\gamma}$	Spatial derivative of A (vector and its elements). In cases where relevant, includes grad-h terms. Can be combined with the above notation that indicates the kernel used
∇A, $\partial^\gamma A$	Gradient of A (vector and its elements). In cases where relevant, does not include grad-h terms
$\nabla_\kappa A$, $\partial^\gamma_\kappa A$	Gradient of A (vector and its elements), interpolated only based on particle neighbours of the same material. In cases where relevant, does not include grad-h terms

APPENDIX B

REMIX SPH equations

$$\frac{d\rho_i}{dt} = \sum_j m_j \, \frac{\rho_i}{\rho_j} \, v_{ij}^{\alpha} \, \frac{1}{2} \left(\left.\frac{d\widetilde{\mathcal{W}}}{dr^{\alpha}}\right|_{ij} - \left.\frac{d\widetilde{\mathcal{W}}}{dr^{\alpha}}\right|_{ji} \right) \; + \; \left(\frac{d\rho_i}{dt}\right)_{\mathrm{difn}} \; + \; \left(\frac{d\rho_i}{dt}\right)_{\mathrm{norm}} \; , \tag{B.0.1}$$

$$\frac{dv_i^{\alpha}}{dt} = -\sum_j m_j \, \frac{P_i + Q_{ij} + P_j + Q_{ji}}{\rho_i \rho_j} \, \frac{1}{2} \left(\left.\frac{d\widetilde{\mathcal{W}}}{dr^{\alpha}}\right|_{ij} - \left.\frac{d\widetilde{\mathcal{W}}}{dr^{\alpha}}\right|_{ji} \right) \; , \tag{B.0.2}$$

$$\frac{du_i}{dt} = \sum_j m_j \, \frac{P_i + Q_{ij}}{\rho_i \rho_j} \, v_{ij}^{\alpha} \, \frac{1}{2} \left(\left.\frac{d\widetilde{\mathcal{W}}}{dr^{\alpha}}\right|_{ij} - \left.\frac{d\widetilde{\mathcal{W}}}{dr^{\alpha}}\right|_{ji} \right) \; + \; \left(\frac{du_i}{dt}\right)_{\mathrm{difn}} \; . \tag{B.0.3}$$

B.1 Kernel gradients

$$\left.\frac{d\widetilde{\mathcal{W}}}{dr^{\gamma}}\right|_{ij} = s_i \left.\frac{d\mathcal{W}}{dr^{\gamma}}\right|_{ij} + (1 - s_i) \left.\frac{dW}{dr^{\gamma}}\right|_{ij} . \tag{B.1.4}$$

$$\left.\frac{d\mathcal{W}}{dr^{\gamma}}\right|_{ij} = A_i B_i^{\alpha} \overline{W}_{ij} + A_i \left(1 + B_i^{\alpha} r_{ij}^{\alpha}\right) \left.\frac{d\overline{W}}{dr^{\gamma}}\right|_{ij} + \left(1 + B_i^{\alpha} r_{ij}^{\alpha}\right) \overline{W}_{ij} \left.\frac{dA}{dr^{\gamma}}\right|_{i} + A_i r_{ij}^{\alpha} \overline{W}_{ij} \left.\frac{dB^{\alpha}}{dr^{\gamma}}\right|_{i} . \tag{B.1.5}$$

B.1.1 Symmetrised kernels and their gradients

$$\overline{W} \equiv \frac{W_{ij} + W_{ji}}{2} \; , \tag{B.1.6}$$

$$\left.\frac{d\overline{W}}{dr^\gamma}\right|_{ij} = \frac{1}{2}\left(\left.\frac{\partial W}{\partial r^\gamma}\right|_{ij} + \left.\frac{\partial W}{\partial h}\right|_{ij} \partial_i^\gamma \hat{h} - \left.\frac{\partial W}{\partial r^\gamma}\right|_{ji} \right) \; , \tag{B.1.7}$$

$$\partial_i^\gamma \hat{h} = \sum_j (h_j - h_i)\, \partial_i^\gamma \hat{W}_{ij} \frac{m_j}{\rho_j} \; , \tag{B.1.8}$$

$$\partial_i^\gamma \hat{W}_{ij} \equiv \frac{\partial_i^\gamma W_{ij}}{m_{0,\,i}} - \frac{W_{ij}}{m_{0,\,i}^2} \partial_i^\gamma m_0 \; . \tag{B.1.9}$$

B.1.2 Linear-order reproducing kernel construction

$$A_i = \left(\overline{m}_{0,\,i} - \left(\overline{m}_{2,\,i}^{-1}\right)^{\alpha\beta} \overline{m}_{1,\,i}^{\alpha}\, \overline{m}_{1,\,i}^{\beta} \right)^{-1} \; , \tag{B.1.10}$$

$$\left.\frac{dA}{dr^\gamma}\right|_i = -A_i^2 \left(\left.\frac{d\overline{m}_0}{dr^\gamma}\right|_i - 2\left(\overline{m}_{2,\,i}^{-1}\right)^{\alpha\beta} \overline{m}_{1,\,i}^{\beta} \left.\frac{d\overline{m}_1^\alpha}{dr^\gamma}\right|_i + \left(\overline{m}_{2,\,i}^{-1}\right)^{\alpha\phi} \left.\frac{d\overline{m}_2^{\phi\psi}}{dr^\gamma}\right|_i \left(\overline{m}_{2,\,i}^{-1}\right)^{\psi\beta} \overline{m}_{1,\,i}^{\alpha} \overline{m}_{1,\,i}^{\beta} \right) \; , \tag{B.1.11}$$

$$B_i^\alpha = -\left(\overline{m}_{2,\,i}^{-1}\right)^{\alpha\beta} \overline{m}_{1,\,i}^{\beta} \; , \tag{B.1.12}$$

$$\left.\frac{dB^\alpha}{dr^\gamma}\right|_i = -\left(\overline{m}_{2,\,i}^{-1}\right)^{\alpha\beta} \left.\frac{d\overline{m}_1^\beta}{dr^\gamma}\right|_i + \left(\overline{m}_{2,\,i}^{-1}\right)^{\alpha\phi} \left.\frac{d\overline{m}_2^{\phi\psi}}{dr^\gamma}\right|_i \left(\overline{m}_{2,\,i}^{-1}\right)^{\psi\beta} \overline{m}_{1,\,i}^{\beta} \; . \tag{B.1.13}$$

$$\overline{m}_{0,\,i} = \sum_j \overline{W}_{ij} V_j \; , \qquad \left.\frac{d\overline{m}_0}{dr^\gamma}\right|_i = \sum_j \left.\frac{d\overline{W}}{dr^\gamma}\right|_{ij} V_j \; , \tag{B.1.14}$$

$$\overline{m}_{1,\,i}^{\alpha} = \sum_j r_{ij}^\alpha \overline{W}_{ij} V_j \; , \qquad \left.\frac{d\overline{m}_1^\alpha}{dr^\gamma}\right|_i = \sum_j \left(r_{ij}^\alpha \left.\frac{d\overline{W}}{dr^\gamma}\right|_{ij} + \delta^{\alpha\gamma} \overline{W}_{ij} \right) V_j \; , \tag{B.1.15}$$

$$\overline{m}_{2,\,i}^{\alpha\beta} = \sum_j r_{ij}^\alpha r_{ij}^\beta \overline{W}_{ij} V_j \; , \quad \left.\frac{d\overline{m}_2^{\alpha\beta}}{dr^\gamma}\right|_i = \sum_j \left(r_{ij}^\alpha r_{ij}^\beta \left.\frac{d\overline{W}}{dr^\gamma}\right|_{ij} + \left(r_{ij}^\alpha \delta^{\beta\gamma} + \delta^{\alpha\gamma} r_{ij}^\beta \right) \overline{W}_{ij} \right) V_j \; . \tag{B.1.16}$$

B.1.3 Vacuum boundary switch

$$s(h_i|\mathbf{B}_i|) = \begin{cases} \exp\left[-\dfrac{\left(0.8 - h_i|\mathbf{B}_i|\right)^2}{0.08} \right] & \text{for } h_i|\mathbf{B}_i| \geq 0.8 \; , \\ 1 & \text{otherwise} \; , \end{cases} \tag{B.1.17}$$

B.2 Artificial viscosity and artificial diffusion

$$Q_{ij} = \frac{1}{2}\left(a_{\text{visc}} + b_{\text{visc}}\mathcal{B}_i^{\text{visc}}\right)\rho_i\left(-\alpha c_i \mu_{ij} + \beta\mu_{ij}^2\right) \ , \tag{B.2.18}$$

with $\alpha = 1.5$, $\beta = 3$, $a_{\text{visc}} = 2/3$, and $b_{\text{visc}} = 1/3$.

$$\mu_{ij} = \begin{cases} \dfrac{\tilde{\mathbf{v}}_{ij}\cdot\boldsymbol{\eta}_{ij}}{\boldsymbol{\eta}_{ij}\cdot\boldsymbol{\eta}_{ij} + \epsilon^2} & \text{for } \tilde{\mathbf{v}}_{ij}\cdot\boldsymbol{\eta}_{ij} < 0 \ , \\ 0 & \text{otherwise} \ , \end{cases} \tag{B.2.19}$$

with $\epsilon = 0.1$ and $\boldsymbol{\eta}_{ij} \equiv (\mathbf{r}_i - \mathbf{r}_j)/h_i$.

$$\left(\frac{du_i}{dt}\right)_{\text{difn}} = \sum_j \kappa_{ij}\left(a_u + b_u\mathcal{B}_{ij}^{\text{difn}}\right) v_{\text{sig},ij}\left(\tilde{u}_j - \tilde{u}_i\right)\frac{m_j}{\rho_{ij}}\frac{1}{2}\left|\left.\frac{d\widetilde{\mathcal{W}}}{d\mathbf{r}}\right|_{ij} - \left.\frac{d\widetilde{\mathcal{W}}}{d\mathbf{r}}\right|_{ji}\right| \ , \tag{B.2.20}$$

$$\left(\frac{d\rho_i}{dt}\right)_{\text{difn}} = \sum_j \kappa_{ij}\left(a_\rho + b_\rho\mathcal{B}_{ij}^{\text{difn}}\right) v_{\text{sig},ij}\left(\tilde{\rho}_j - \tilde{\rho}_i\right)\frac{\rho_i}{\rho_j}\frac{m_j}{\rho_{ij}}\frac{1}{2}\left|\left.\frac{d\widetilde{\mathcal{W}}}{d\mathbf{r}}\right|_{ij} - \left.\frac{d\widetilde{\mathcal{W}}}{d\mathbf{r}}\right|_{ji}\right| \ , \tag{B.2.21}$$

with $a_u = a_\rho = 0.05$, $b_u = b_\rho = 0.95$, $v_{\text{sig},\,ij} = |\tilde{\mathbf{v}}_i - \tilde{\mathbf{v}}_j|$, $\rho_{ij} \equiv (\rho_i + \rho_j)/2$, and $\kappa_{ij} = 1$ for particles of the same material and $\kappa_{ij} = 0$ otherwise.

B.2.1 Quantities reconstructed to particle midpoints

$$\tilde{v}_{ij}^{\alpha} = v_i^{\alpha} + \frac{1}{2}\left(1 - \mathcal{B}_i^{\text{SL}}\right)\Phi_{v,\,ij}\left(r_j^{\gamma} - r_i^{\gamma}\right)\partial_i^{\gamma}\hat{v}^{\alpha} \ , \quad \partial_i^{\gamma}\hat{v}^{\alpha} = \sum_j \left(v_j^{\alpha} - v_i^{\alpha}\right)\partial_i^{\gamma}\hat{W}_{ij}\frac{m_j}{\rho_j} \ , \tag{B.2.22}$$

$$\tilde{u}_i = u_i + \frac{1}{2}\Phi_{u,\,ij}\left(r_j^{\gamma} - r_i^{\gamma}\right)\partial_{\kappa,\,i}^{\gamma}\,\hat{u} \ , \qquad \partial_{\kappa,\,i}^{\gamma}\,\hat{u} = \sum_j \kappa_{ij}\left(u_j - u_i\right)\partial_i^{\gamma}\hat{W}_{ij}\frac{m_j}{\rho_j} \ , \tag{B.2.23}$$

$$\tilde{\rho}_i = \rho_i + \frac{1}{2}\Phi_{\rho,\,ij}\left(r_j^{\gamma} - r_i^{\gamma}\right)\partial_{\kappa,\,i}^{\gamma}\,\hat{\rho} \ , \qquad \partial_{\kappa,\,i}^{\gamma}\,\hat{\rho} = \sum_j \kappa_{ij}\left(\rho_j - \rho_i\right)\partial_i^{\gamma}\hat{W}_{ij}\frac{m_j}{\rho_j} \ . \tag{B.2.24}$$

B.2.2 Slope limiter

$$\Phi_{ij} = \begin{cases} 0 & \text{for } A_{ij} < 0 \, , \\ \dfrac{4A_{ij}}{(1+A_{ij})^2} \exp\left[- \left(\dfrac{\eta_{ij}^{\min} - \eta_{\text{crit}}}{0.2} \right)^2 \right] & \text{for } \eta_{ij}^{\min} < \eta_{\text{crit}} \, , \\ \dfrac{4A_{ij}}{(1+A_{ij})^2} & \text{otherwise} \, , \end{cases} \tag{B.2.25}$$

$$\eta_{\text{crit}} = \frac{1}{h_i} \left(\frac{1}{\sum_j W_{ij}} \right)^{1/d} \equiv \frac{1}{\eta_{\text{kernel}}} \, , \tag{B.2.26}$$

$$A_{v,\,ij} = \frac{\partial_i^\beta \hat{v}^\alpha (\mathbf{r}_j - \mathbf{r}_i)^\beta (\mathbf{r}_j - \mathbf{r}_i)^\alpha}{\partial_j^\gamma \hat{v}^\phi (\mathbf{r}_j - \mathbf{r}_i)^\gamma (\mathbf{r}_j - \mathbf{r}_i)^\phi} \, , \quad A_{u,\,ij} = \frac{\partial_{\kappa,\,i}^\alpha \, \hat{u} (\mathbf{r}_j - \mathbf{r}_i)^\alpha}{\partial_{\kappa,\,j}^\beta \, \hat{u} (\mathbf{r}_j - \mathbf{r}_i)^\beta} \, , \qquad A_{\rho,\,ij} = \frac{\partial_{\kappa,\,i}^\alpha \, \hat{\rho} (\mathbf{r}_j - \mathbf{r}_i)^\alpha}{\partial_{\kappa,\,j}^\beta \, \hat{\rho} (\mathbf{r}_j - \mathbf{r}_i)^\beta} \, . \tag{B.2.27}$$

B.2.3 Balsara switch

$$\mathcal{B}_i^{\text{visc}} \equiv \mathcal{B}_i^{\text{SL}} \equiv \mathcal{B}_i = \frac{|\nabla \cdot \mathbf{v}_i|}{|\nabla \cdot \mathbf{v}_i| + |\nabla \times \mathbf{v}_i| + 0.0001 c_i / h_i} \, , \qquad \mathcal{B}_{ij}^{\text{difn}} = \frac{\mathcal{B}_i + \mathcal{B}_j}{2} \, . \tag{B.2.28}$$

B.3 Normalising term

$$\left(\frac{d\rho_i}{dt} \right)_{\text{norm}} = \alpha_{\text{norm}} \, s_i \, (m_{0,\,i} - 1) \, \rho_i \sum_j v_{\text{norm},\,ij} \, \frac{m_j}{\rho_{ij}} \frac{1}{2} \left| \left. \frac{d\widetilde{\mathcal{W}}}{d\mathbf{r}} \right|_{ij} - \left. \frac{d\widetilde{\mathcal{W}}}{d\mathbf{r}} \right|_{ji} \right| \, , \tag{B.3.29}$$

with $\alpha_{\text{norm}} = 1,\ v_{\text{norm},\,ij} = |\mathbf{v}_i - \mathbf{v}_j| \, ,\ \rho_{ij} \equiv (\rho_i + \rho_j)/2,$ and $m_{0,\,i} = \sum_j W_{ij} V_j$.

APPENDIX C

Traditional SPH formulations used for comparison simulations

$$\langle \rho_i \rangle = \sum_j m_j W_{ij} \ , \tag{C.0.1}$$

$$\frac{dv_i^\alpha}{dt} = -\sum_j m_j \left(\frac{f_{ij}\, P_i}{\langle \rho_i \rangle^2} \left.\frac{\partial W}{\partial r^\alpha}\right|_{ij} - \frac{f_{ji}\, P_j}{\langle \rho_j \rangle^2} \left.\frac{\partial W}{\partial r^\alpha}\right|_{ji} \right) + \left(\frac{dv_i^\alpha}{dt} \right)_{\rm visc} \ , \tag{C.0.2}$$

$$\frac{du_i}{dt} = \sum_j m_j \frac{f_{ij}\, P_i}{\langle \rho_i \rangle^2} v_{ij}^\alpha \left.\frac{\partial W}{\partial r^\alpha}\right|_{ij} + \left(\frac{du_i}{dt} \right)_{\rm visc} \ . \tag{C.0.3}$$

C.1 Gradient of smoothing length factor

$$\langle n_i \rangle = \sum_j W_{ij} \ , \tag{C.1.4}$$

$$f_{ij} = 1 - \frac{1}{m_j} \frac{h_i}{d \, \langle \rho_i \rangle} \left.\frac{\partial \langle \rho \rangle}{\partial h}\right|_i \left(1 + \frac{h_i}{d \, \langle n_i \rangle} \left.\frac{\partial \langle n \rangle}{\partial h}\right|_i \right)^{-1} \ , \tag{C.1.5}$$

where d is the number of spatial dimensions.

C.2 Artificial viscosity

$$\left(\frac{dv_i^\alpha}{dt}\right)_{\rm visc} = -\sum_j \frac{\mathcal{B}_i+\mathcal{B}_j}{2}\left(-\alpha c_{ij}\,\mu_{ij}+\beta\mu_{ij}^2\right)\frac{m_j}{\langle\rho_{ij}\rangle}\,\frac{1}{2}\left(f_{ij}\left.\frac{\partial W}{\partial r^\alpha}\right|_{ij} - f_{ji}\left.\frac{\partial W}{\partial r^\alpha}\right|_{ji}\right)\,, \tag{C.2.6}$$

$$\left(\frac{du_i}{dt}\right)_{\rm visc} = \frac{1}{2}\sum_j \frac{\mathcal{B}_i+\mathcal{B}_j}{2}\left(-\alpha c_{ij}\,\mu_{ij}+\beta\mu_{ij}^2\right)\frac{m_j}{\langle\rho_{ij}\rangle}\,f_{ij}\,v_{ij}^\alpha\left.\frac{\partial W}{\partial r^\alpha}\right|_{ij}\,, \tag{C.2.7}$$

$$\mu_{ij} = \begin{cases} \dfrac{\mathbf{v}_{ij}\cdot\mathbf{r}_{ij}}{|\mathbf{r}_{ij}|} & \text{for } \mathbf{v}_{ij}\cdot\mathbf{r}_{ij} < 0\,, \\ 0 & \text{otherwise}\,, \end{cases} \tag{C.2.8}$$

with $\alpha = 1.5$, $\beta = 3$, $\langle\rho_{ij}\rangle \equiv (\langle\rho_i\rangle + \langle\rho_j\rangle)/2$, $c_{ij} \equiv (c_i + c_j)/2$, $\mathbf{v}_{ij} \equiv \mathbf{v}_i - \mathbf{v}_j$, $\mathbf{r}_{ij} \equiv \mathbf{r}_i - \mathbf{r}_j$, and

$$\mathcal{B}_i = \frac{|\nabla\cdot\mathbf{v}_i|}{|\nabla\cdot\mathbf{v}_i| + |\nabla\times\mathbf{v}_i| + 0.0001 c_i/h_i}\,. \tag{C.2.9}$$

C.3 Artificial conduction

Only used where indicated. Based on the SPHENIX implementation (Borrow et al., 2022).

$$\frac{du_i}{dt} = \frac{P_i}{\langle\rho_i\rangle^2}\sum_j m_j\,v_{ij}^\alpha\left.\frac{\partial W}{\partial r^\alpha}\right|_{ij} + \left(\frac{du_i}{dt}\right)_{\rm visc} + \left(\frac{du_i}{dt}\right)_{\rm cond}\,, \tag{C.3.10}$$

$$\left(\frac{du_i}{dt}\right)_{\rm cond} = \sum_j \alpha_{c,\,ij}\,v_{c,\,ij}\,(u_i - u_j)\,\frac{m_j}{\langle\rho_{ij}\rangle}\,\frac{1}{2}\left(f_{ij}\left.\frac{\partial W}{\partial r^\alpha}\right|_{ij} - f_{ji}\left.\frac{\partial W}{\partial r^\alpha}\right|_{ji}\right)\,, \tag{C.3.11}$$

$$\alpha_{c,\,ij} = \frac{P_i\,\alpha_{c,\,i} + P_j\,\alpha_{c,\,j}}{P_i + P_j}\,, \qquad v_{c,\,ij} = \frac{1}{2}\left(\frac{|\mathbf{v}_{ij}\cdot\mathbf{r}_{ij}|}{|\mathbf{r}_{ij}|} + \sqrt{\frac{|P_i - P_j|}{\langle\rho_{ij}\rangle}}\right)\,, \tag{C.3.12}$$

where $\alpha_{c,\,i}$ is evolved in time as described by Borrow et al. (2022).

APPENDIX D

Kernel choice

Here we present results from square test (§2.3.1), Sod shock tube (§2.3.2), ideal gas KHI (§2.3.3), and KHI with Earth-like iron & rock (§2.3.4) simulations to show the effect of smoothing kernel choice in REMIX simulations. We use five different kernel functions with corresponding η_{kernel}: cubic spline with $\eta_{\text{kernel}} = 1.292$ (∼55 neighbours); quartic spline with $\eta_{\text{kernel}} = 1.203$ (∼60 neighbours); Wendland C^2 with $\eta_{\text{kernel}} = 1.487$ (∼100 neighbours); Wendland C^4 with $\eta_{\text{kernel}} = 1.643$ (∼200 neighbours); Wendland C^6 with $\eta_{\text{kernel}} = 1.866$ (∼400 neighbours) (Dehnen & Aly, 2012).

Square tests, presented in Fig. D.1, show similar behaviour in the equal spacing scenario for all of these kernels, although with more noise in the lower-order kernels with fewer neighbours. For the higher-order Wendland kernels, results are very similar over these timescales, with very little particle motion. In the equal mass scenario, however, the use of either the Wendland C^4 or Wendland C^6 kernels leads to spurious behaviour at the corners. The particle noise in simulations with lower-order kernels is in fact helpful in disturbing the growth of these slowly evolving, error-driven features. We found that using grad-h terms calculated directly from Eqn. 1.4 combined with a higher artificial viscosity helps the treatment of corners

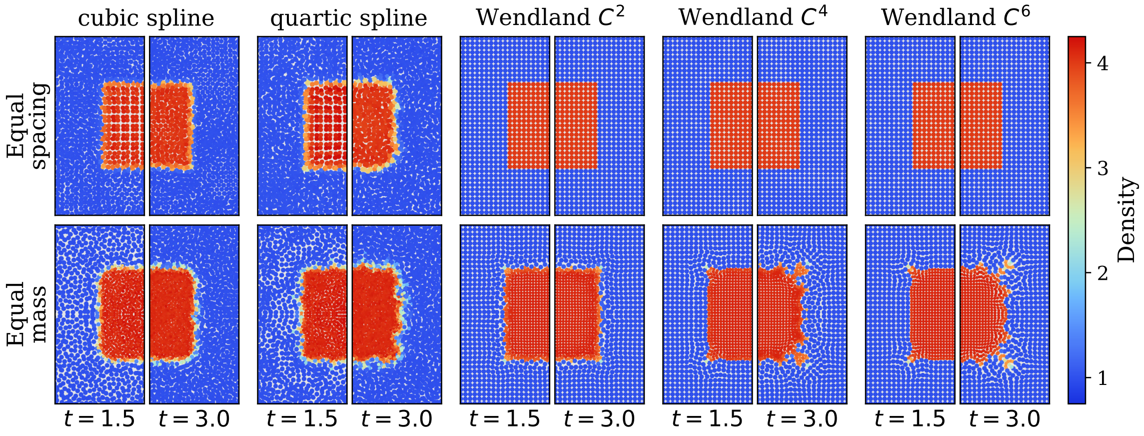

Figure D.1: The effect of kernel choice in square test simulations. We plot snapshots from REMIX square test simulations with equal initial particle spacing and equal particle mass, for 5 different kernel functions at two times. Particles are coloured by their density.

in the square tests with higher-order kernels. However, zeroth-order error in grad-h terms calculated in that way leads to problematic behaviour in regions away from density discontinuities, and we choose to take a conservative artificial viscosity approach, keeping it low away from shocks. Therefore, we choose to calculate grad-h terms as we describe in §2.2.3. Again, lower-order kernels show more particle noise, however, the cross-section of the cube does not lose its square shape.

In the Sod shock (Fig. D.2) and both KHI tests (Figs. D.3 and D.4) We see a general trend of higher-order kernels reducing particle noise. However, these effects are minor compared with the primary improvements in all these simulations compared with traditional SPH equivalents.

Based on these simulations, we conclude that the Wendland C^2 kernel is a good compromise between accuracy and computational speed, which is why we use it for all simulations other than those presented in this section. Using a higher-order kernel only leads to small improvements in noise reduction in these tests and, in the case of the square test, gives worse results. We stress, however, that the behaviour at the corners of a 3D cube is not necessarily important for many science applications, if other benefits are offered in more typical configurations. The lower-order kernels lead to more particle noise, however, these still show significant improvements compared with traditional SPH simulations. This suggests that, for example a cubic spline kernel could be used with REMIX for applications where simulation run-speed is an

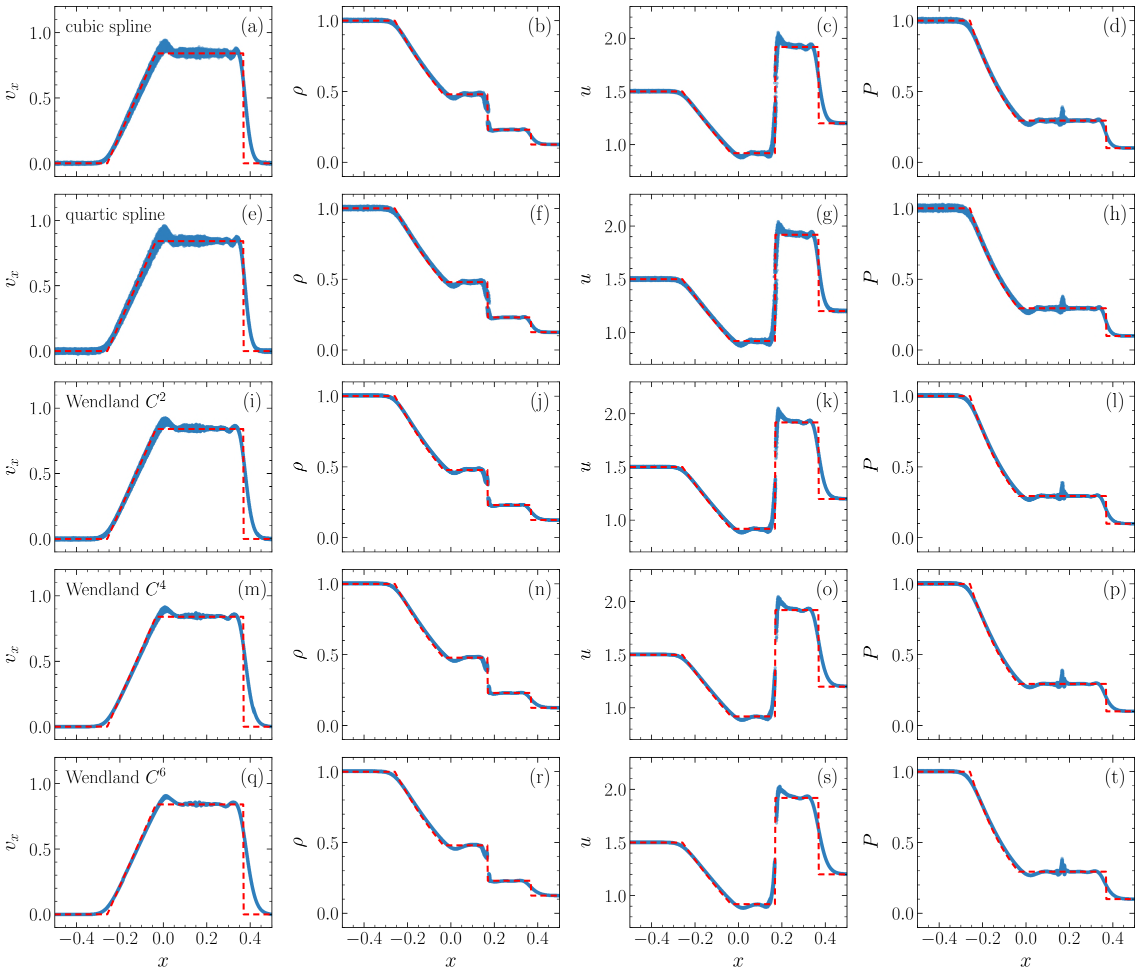

Figure D.2: The effect of kernel choice in REMIX Sod shock tube simulations. Velocity in the x-direction, v_x, density, ρ, specific internal energy, u, and pressure, P plotted against x-position at time $t = 0.2$. Rows correspond to the 5 different kernel functions. The reference solution is plotted in red.

important consideration.

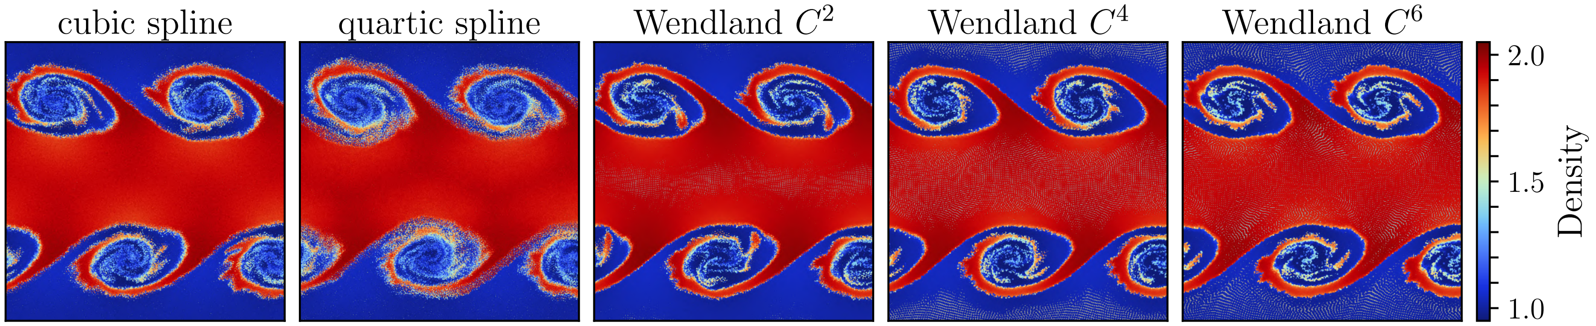

Figure D.3: The effect of kernel choice in REMIX ideal gas KHI simulation. Plotted at time $t = 2\ \tau_{\mathrm{KH}}$, for 5 different kernel functions. Particles are coloured by their density.

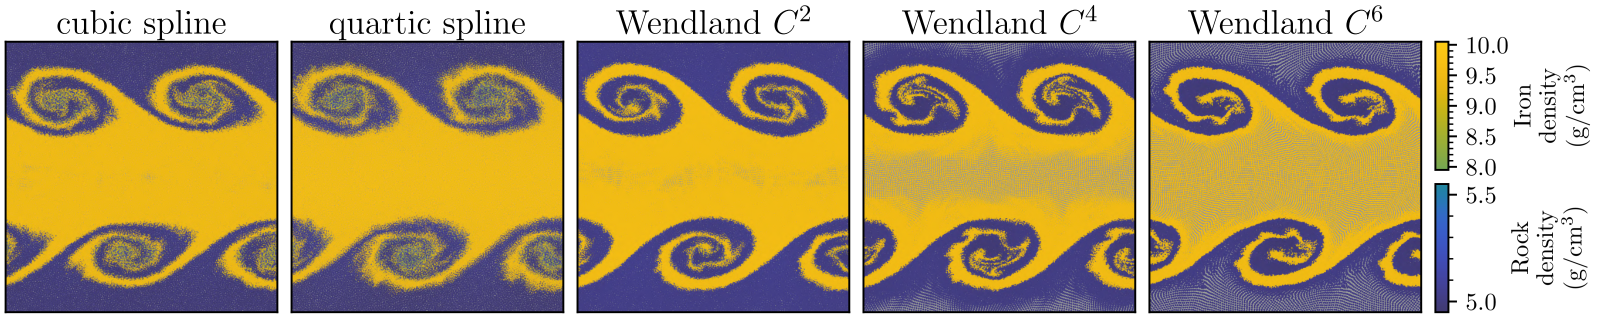

Figure D.4: The effect of kernel choice in REMIX KHI simulations between Earth-like iron & rock. Plotted at time $t = 2\ \tau_{\mathrm{KH}}$, for 5 different kernel functions. Particles are coloured by their material type and density.

APPENDIX E

Choices made in linear-order reproducing kernel construction

In Fig. E.1, sensitivities to subtle choices in the kernel construction (§2.2.3) are presented in the context of the square test with equal mass particles (§2.3.1). We include panels to show equivalent cases without artificial diffusion of density or internal energy in the model. This isolates the effect of kernel choice from stabilising effects caused by stronger diffusion for particles with higher relative speeds.

We consider using W_{ij} rather than the averaged kernel $\overline{W}_{ij} \equiv [W(\mathbf{r}_{ij}, h_i) +$

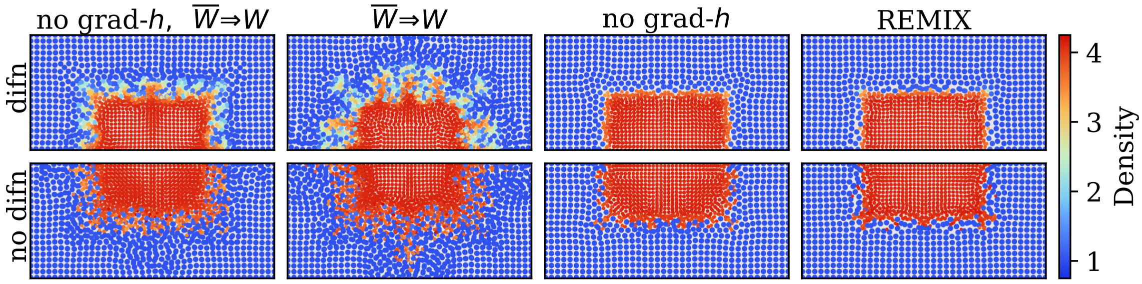

Figure E.1: The effect of choices made in the construction of linear-order reproducing kernels in square test simulations. Plots show a central cross-section from square test simulations with equal mass particles, using subtly different variations of REMIX. Snapshots are at time $t = 3.0$. We show the results with (top) and without (bottom) diffusion of density and internal energy, for the full REMIX scheme and with combinations of either no grad-h terms and without using an averaged kernel in the construction of our linear-order reproducing kernels.

$W(\mathbf{r}_{ji}, h_j)]/2$. This permeates through the full reproducing kernel construction, including geometric moments and their gradients. Switching to W_{ij} also increases grad-h terms by a factor of 2, as seen in Eqn. 2.17. Using $\overline{W}_{ij}$ in the construction of the reproducing kernels is shown to be important in achieving good behaviour in the square test, with all W_{ij} cases showing a disruption of the cube by particle motions at the interface. This demonstrates how additional error introduced in kernel antisymmetrisation (for conservation) in the equations of motion is sensitive to the base-kernel used in the construction of the reproducing kernels. Regardless of the kernel used in the construction, the cube would remain undisturbed if the antisymmetrisation step was not carried out.

The inclusion of grad-h terms only leads to a small effect. We note again that we include these terms primarily for completeness of the methods, and since they have no negative impact on run speed or any other considerations.

APPENDIX F

Choices made in artificial viscosity and diffusion construction

Here we show how Sod shock (§2.3.2) and ideal gas Kelvin–Helmholtz instability (§2.3.3) simulations motivate the parameters used in the REMIX artificial viscosity and artificial diffusion schemes. We (1) demonstrate the effectiveness of a Balsara switch in the viscosity slope limiter in reducing oscillations in shocks; (2) motivate our choice of the standard viscosity constants, α and β; (3) motivate our choice of the parameters that distinguish between the treatment of artificial viscosity in shocks and shearing regions, $a_{\rm visc}$ and $b_{\rm visc}$; (4) motivate our choice of the parameters that distinguish between the treatment of artificial diffusion in shocks and shearing region, $a_{\rm difn} \equiv a_u = a_\rho$ and $b_{\rm difn} \equiv b_u = b_\rho$. For simplicity, we use the same parameters for the diffusion of internal energy and density. Our approach to these artificial terms is a conservative one, in which we deliberately keep them weak, while still strong enough to give noticeable improvements.

The Sod shock and KHI examples were chosen since the artificial viscosity and diffusion play different roles in these two scenarios: we require strong viscosity and diffusion to accurately capture shocks, whereas in shearing regions they are only required to smooth out accumulated noise in particle velocities, densities, and internal energies. To isolate these elements of the construction we use simplified

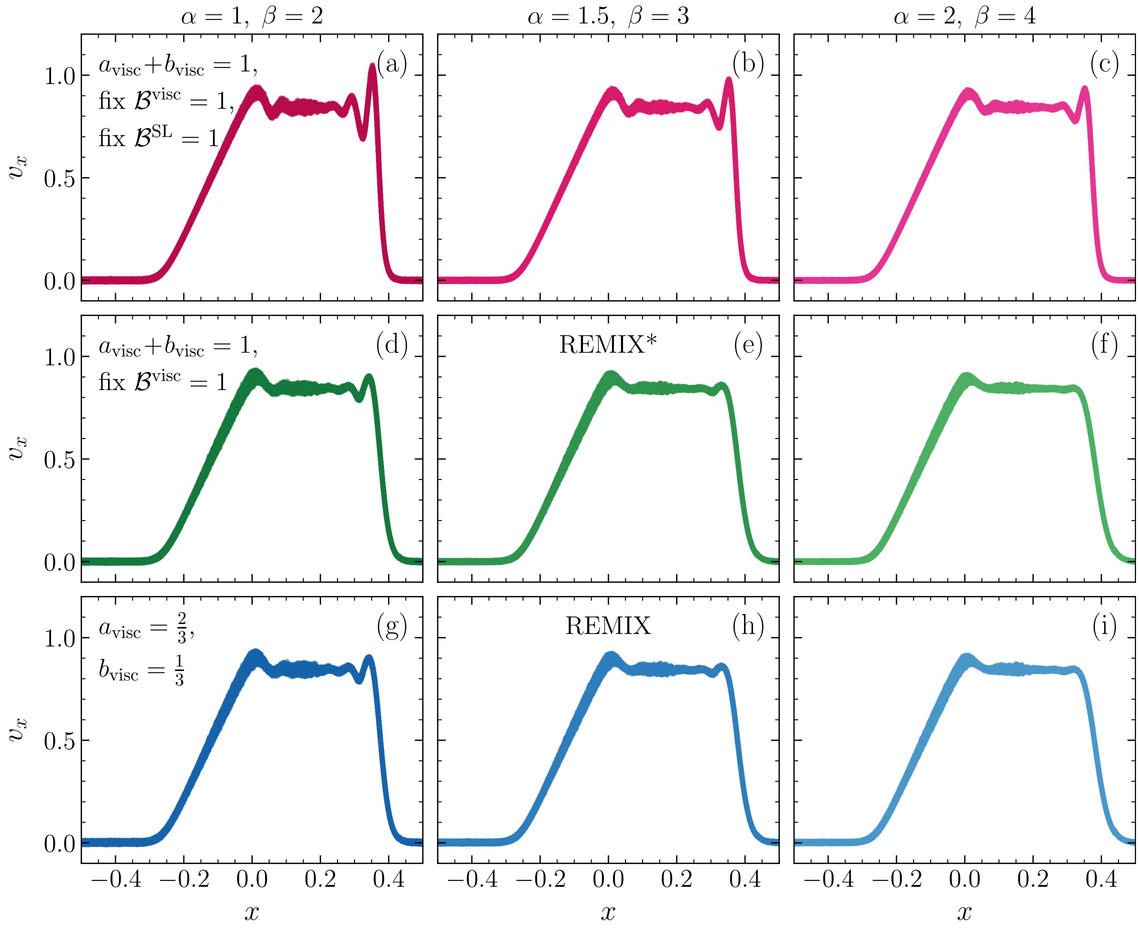

Figure F.1: The effect of artificial viscosity parameters in Sod shock tube simulations. Velocity in the x-direction, v_x, plotted against x-position from Sod shock tube simulations at time $t = 0.2$, using REMIX with variations in artificial viscosity formulation. The annotations "REMIX" and "REMIX*" correspond to panels showing the final REMIX scheme and the equivalent simplified shock case respectively.

versions of the REMIX construction, in which we overwrite the Balsara switches that appear in Eqns. 2.29–2.31, to 1 or 0 to isolate the treatment in shocks or shearing regions respectively: In the simplified shock case, we set $\mathcal{B} = 1$, such that the strength of viscosity or diffusion is parameterised by the respective $a + b$ in Eqns. 2.29–2.31. In the simplified shear case, we set $\mathcal{B} = 0$, such that the strengths are set by a only. Therefore, we use these to select values for a and b parameters, based on which the full scheme can then switch smoothly between these two simplified versions.

We first focus on the artificial viscosity scheme. In Fig. F.1 we investigate the effect of changes in the viscosity treatment on Sod shock tube simulations. Rows correspond to: the simplified shock case with $a_{\rm visc} + b_{\rm visc} = 1$ and Balsara switches

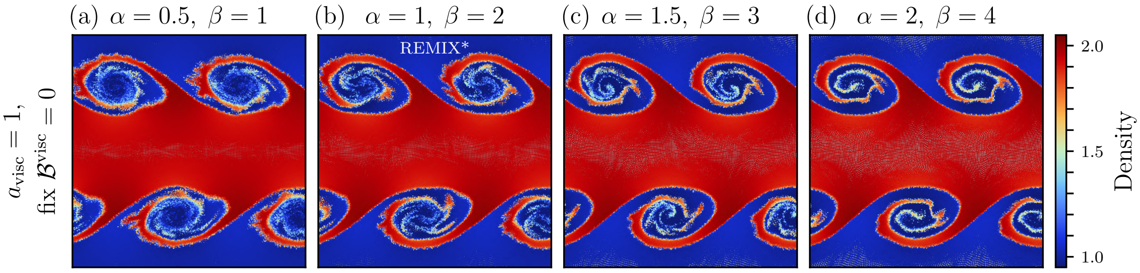

Figure F.2: The effect of artificial viscosity parameters in ideal gas KHI simulations with $\mathcal{B}^{\rm visc}$ fixed to 0. Snapshots are from simulations with sharp initial discontinuities, at time $t = 2\ \tau_{\rm KH}$, using REMIX with variations in artificial viscosity formulation. Particles are coloured by their density. The annotation “REMIX*” corresponds to the panel showing the simplified shear case that is equivalent to the final REMIX scheme.

in the viscosity slope limiter set to $\mathcal{B}^{\rm SL} = 1$; the simplified shock case again with $a_{\rm visc} + b_{\rm visc} = 1$, but without fixing $\mathcal{B}^{\rm SL}$; and the final REMIX scheme. Columns correspond to different choices of α and β, with a consistent $\beta = 2\alpha$ in all cases here. We note that fixing this ratio leads to a degeneracy for α and β with $a_{\rm visc} + b_{\rm visc}$ in the simplified shock case, and therefore we choose to set $a_{\rm visc} + b_{\rm visc} = 1$.

We first note that the Balsara switch $\mathcal{B}^{\rm SL}$ is effective in dissipating oscillations in the shock. The oscillations in the case that makes use of $\mathcal{B}^{\rm SL}$ but has the weakest viscosity (Fig. F.1(d)) are smaller than even those in F.1(c), where the α and β factors are twice as large, but $\mathcal{B}^{\rm SL} = 1$ is fixed. Using $\mathcal{B}^{\rm SL}$ therefore allows us to target viscosity to shocks by more effectively switching off the linear reconstruction of velocities to particle midpoints. This in turn allows us to reduce the viscosity parameters so that we can construct a less dissipative artificial viscosity scheme. We note that Fig. F.1(c) is equivalent to the viscosity construction of Frontiere et al. (2017) and Rosswog (2020), other than the latter’s use of quadratic reconstruction, which only makes a minor difference.

Next we consider how the Sod shock is sensitive to the primary viscosity parameters α and β for $a_{\rm visc} + b_{\rm visc} = 1$, as shown in Fig. F.1(d)–(f). Increasing α and β in this way uniformly increases the strength of the viscosity. With $\alpha = 2$ and $\beta = 4$, oscillations are effectively removed. With $\alpha = 1$ and $\beta = 2$, oscillations in the shock are still large compared with the particle scatter. $\alpha = 1.5$ and $\beta = 3$, which we eventually use in REMIX, still has some oscillations, but these are small.

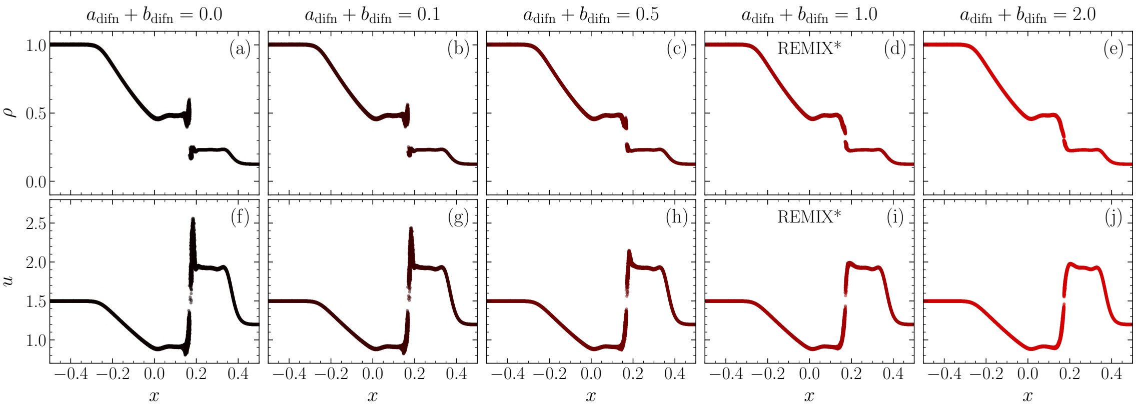

Figure F.3: The effect of artificial diffusion parameters in Sod shock tube simulations with $\mathcal{B}^{\rm difn}$ fixed to 1. Density, ρ, and specific internal energy, u, plotted against x-position at time $t = 0.2$, using REMIX with variations in artificial diffusion formulation. The annotation "REMIX*" corresponds to the panels showing the simplified shock case that is equivalent to the final REMIX scheme.

The bottom row of figures demonstrates that our simplification of this shock case with $\mathcal{B}^{\rm visc} = 1$ is appropriate as we see little difference when comparing with a full REMIX-like construction.

We use the ideal gas KHI in the simplified shear case, combined with the above Sod shock results, to make decisions for values of the REMIX viscosity parameters. In Fig. F.2 we see the effect of changing α and β with $a_{\rm visc} = 1$ and $\mathcal{B}^{\rm visc}$ fixed to 0. Again we note the degeneracy in increasing α and β, this time with increasing $a_{\rm visc}$. With increased artificial viscosity, the boundary of the spiralling KHI plume becomes more pronounced, with less mixing of particles across the interface. In the lowest viscosity case, the small-scale spirals are diffused and structure is not maintained.

For the REMIX scheme we take a conservative approach and choose a viscosity model that, with the assumptions of the simplified approaches considered here, switches between: F.1(e) in shocks, in which oscillations have mostly, but not fully, been removed; to Fig. F.2(b), in which small-scale KHI structure persists, but mixing on the particle scale is not strongly suppressed by artificial viscosity. We note that, in practice, the artificial viscosity will be slightly between the two cases we aim to switch between. This was kept in mind when making this choice. This corresponds to choices of $\alpha = 1.5$, $\beta = 3$ and $a_{\rm visc} = 2/3$, $b_{\rm visc} = 1/3$.

Next we consider the artificial diffusion model. We use the same parameter

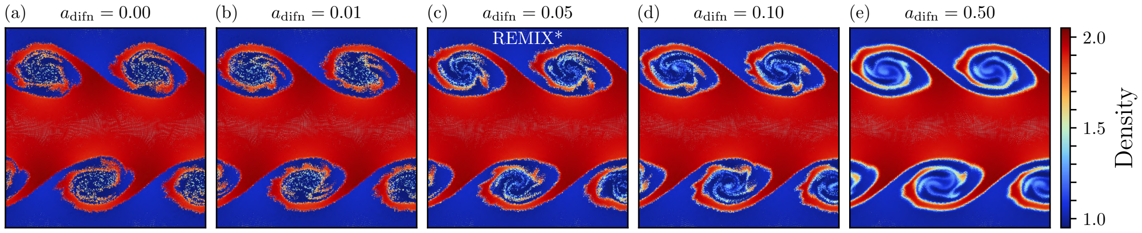

Figure F.4: The effect of artificial diffusion parameters in ideal gas KHI simulations with $\mathcal{B}^{\rm difn}$ fixed to 0. Snapshots are from simulations with sharp initial discontinuities, at time $t = 2\ \tau_{\rm KH}$, using REMIX with variations in artificial diffusion formulation. Particles are coloured by their densities. The annotation "REMIX*" corresponds to the panel showing the simplified shear case that is equivalent to the final REMIX scheme.

values in the artificial diffusion of density and internal energy, for simplicity. First we consider the Sod shock in the simplified shock case, however this time applied to the diffusion equations (Eqns. 2.30 and 2.31). In Fig. F.3, we demonstrate our need for artificial diffusion of both density and internal energy: in the first column, we see the case of no diffusion leading to sizable spikes in these quantities. As the strength of artificial diffusion is increased, the spikes are smoothed.

In Fig. F.4, we show snapshots from ideal gas KHI simulations in the simplified shear case. At high diffusion, the density discontinuity is smoothed and a sharp interface is not maintained. However, some diffusion is helpful in stabilising the evolution of the instability, as we see that with little or no diffusion the structure in the inner regions of the vortex is dominated by particle noise.

Again, for the REMIX scheme we take a conservative approach and choose a diffusion model that, with the assumptions of the simplified approaches considered here, switches between: F.3(d) and (i) for shocks, where spikes in density and internal energy have mostly been removed; to Fig. F.4(c), in which diffusion helps to stabilise the small-scale structure of the instability, but density discontinuities are allowed to persist. We note that, in practice, the artificial diffusion will be slightly between the two cases we aim to switch between. This was kept in mind when making this choice. This corresponds to a choice of $a_{\rm difn} = 0.05$, $b_{\rm difn} = 0.95$. We also note that this can be seen as switching from a weak diffusion away from shocks with a factor of 0.05, similar to that of Rosswog (2020), to a stronger diffusion in shocks with a factor of 1, similar to Price et al. (2018).

APPENDIX G

Further Kelvin–Helmholtz results and figures

As discussed in §2.3.3, a discontinuous shearing interface will be unstable to the growth of perturbations of all wavenumbers (Robertson et al., 2010). In simulations of Kelvin–Helmholtz instabilities, the wavenumbers of modes that are allowed to grow is limited by the numerical resolution; modes of wavelengths of the particle-separation scale and shorter will not be resolved. Here we demonstrate how error-seeded secondary modes will inevitably grow in KHI simulations with higher-resolution REMIX. We note that although this means that it is impossible to reach a converged solution in this scenario such that we cannot quantitatively judge the accuracy of these results, the presence of secondary modes is still a positive sign that spurious surface tension-like effects are not dominating behaviour at density discontinuities.

First, we consider ideal gas KHI simulations (§2.3.3). In Fig. G.1, we show snapshots from tSPH, tSPH with conduction, and REMIX simulations of KHIs at different resolutions. Additionally, in the bottom row of panels, we show how the shearing interface evolves with REMIX if no initial velocity perturbation is applied to the system. We see strong surface tension-like effects in the tSPH simulations for all resolutions. Artificial conduction is helpful as resolution is increased, however,

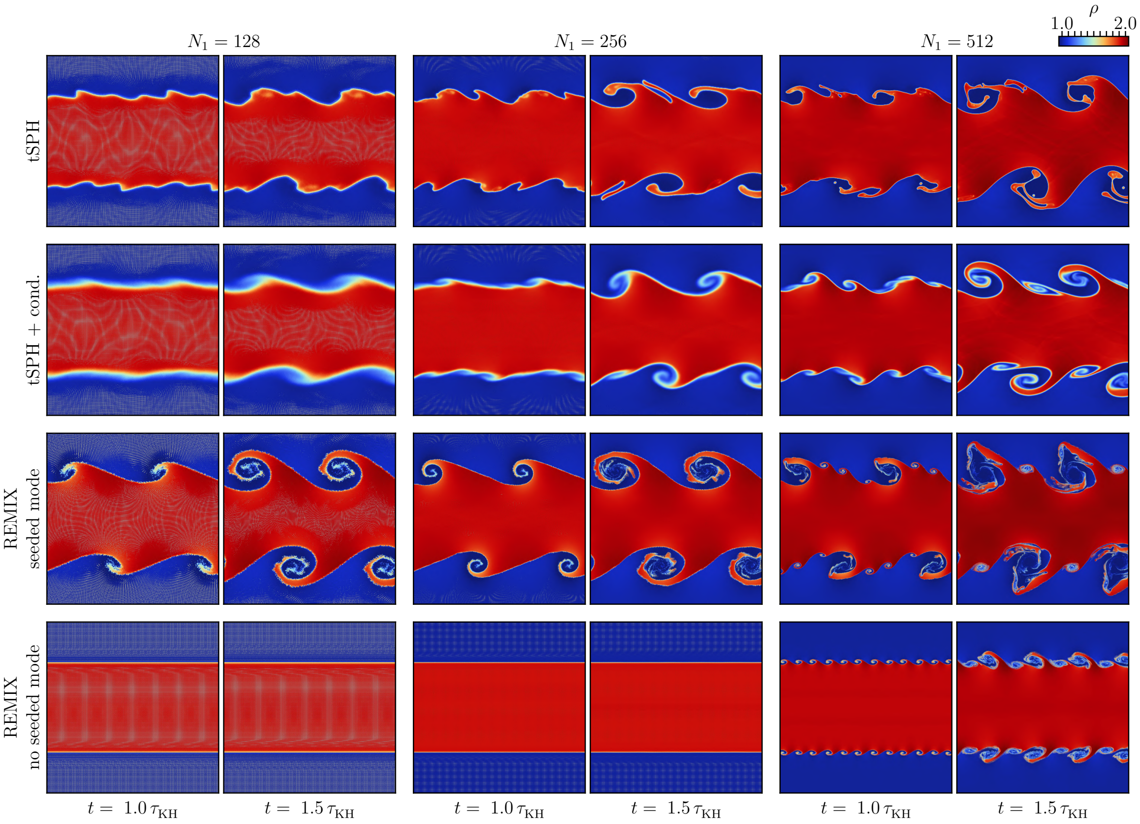

Figure G.1: The effect of resolution in ideal gas KHI simulations with a sharp discontinuity, carried out using tSPH, tSPH with conduction and REMIX. We also plot REMIX simulations without a deliberately seeded perturbation. The KHI is plotted at two times for simulations of three resolutions. Particles are coloured by their density.

the growth of the instability is still slow and the discontinuity becomes diffuse. Over these timescales, the $N_1 = 128$ and $N_1 = 256$ REMIX simulations are largely undisturbed by secondary modes, both in the cases with and without the seeded mode. In contrast, in the higher-resolution $N_1 = 512$ case, we see that secondary modes grow to greatly affect the evolution of the system, both with and without a deliberately seeded mode. The secondary modes grow over the same timescale in both these cases, demonstrating that these are purely seeded by error and noise in the numerical methods, rather than being associated with the growth of the primary mode. We note that over longer timescales than $\tau_{\rm KH}$, secondary modes will also grow in the lower-resolution simulations, although they do not greatly influence the early growth of the instability.

In Fig. G.2, we show similar snapshots from Earth-like KHI simulations at differ-

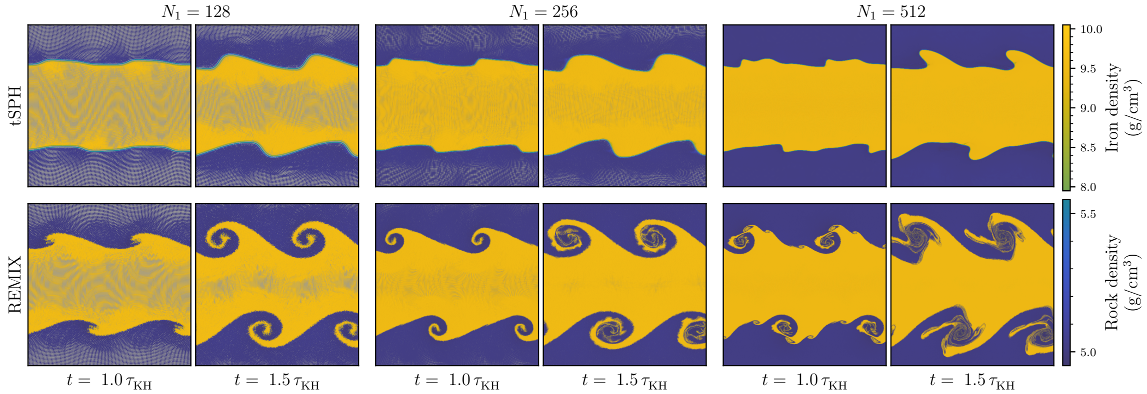

Figure G.2: The effect of resolution in Earth-like iron & rock KHI simulations. Snapshots show two times for simulations of three resolutions, carried out using tSPH and REMIX. Particles are coloured by their material type and density.

ent resolutions with tSPH and REMIX, constructed equivalently to those presented in §2.3.4. Here surface tension-like effects are very strong in tSPH simulations at all resolutions. REMIX is able to deal with this challenging scenario, even at the lowest resolution simulated. Again, here we see error-seeded secondary modes in the $N_1 = 512$ REMIX simulation that grow to length scales where they play a significant role in the evolution of the system. In the lower-resolution simulations, these modes are not resolved and so the primary mode is undisturbed over these timescales.

APPENDIX H

Further planetary results and figures

Here we present further results regarding the planetary settling examples (§2.3.9) to expand on the issues faced in these simulations. First, in Fig. H.1, we show radial profiles, equivalent to those presented in Fig. 2.17, but at the start of the simulation rather than at a time when particles have evolved from their initial configuration. Here we notice in the tSPH case that the blips in pressure are even larger, and that there is also noticeable smoothing of the density field at the vacuum boundary in the Earth-like planet. At the later times, plotted in Fig. 2.17, particles have evolved to

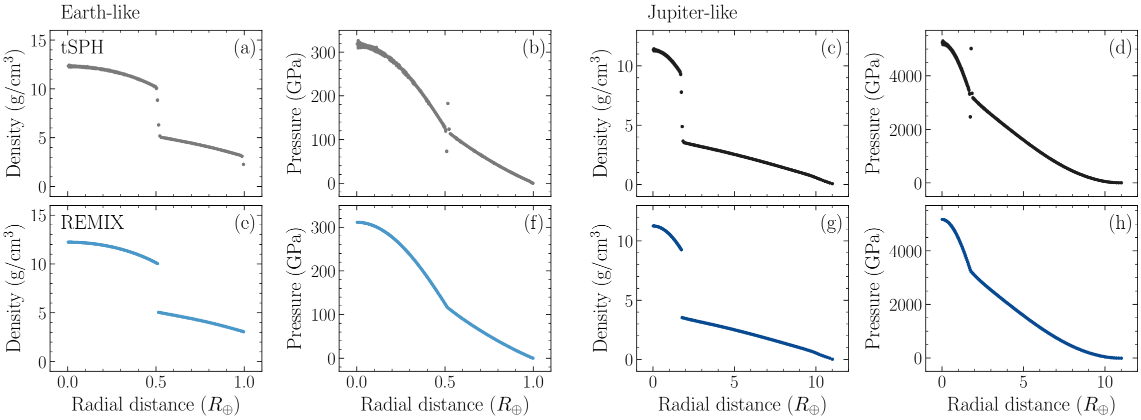

Figure H.1: Radial profiles of density and pressure for Earth-like (a, b, e, f) and Jupiter-like (c, d, g, h) planets at the initial time $t = 0$. Plots show profiles from simulations using tSPH (a–d) and REMIX (e–h).

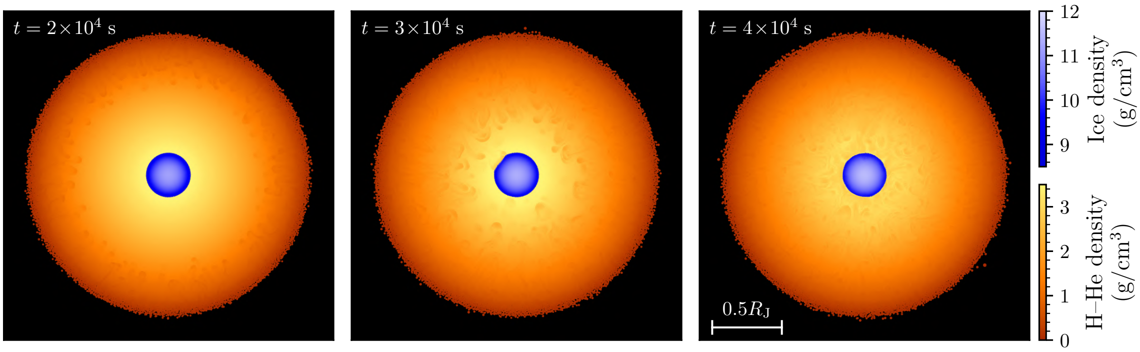

Figure H.2: Snapshots showing the evolution of spurious instabilities in a Jupiter-like planet from a simulation without the density evolution normalising term.

closer to equilibrium configurations where these issues do not appear to be extreme. However, we note that in a more kinematically interesting simulation with particles approaching the material interface, they will encounter the erroneous pressures that lead to surface-tension like effects. Therefore, it is not the case that as the planets relax, they reach a state where spurious surface tension-like forces disappear. We note that both material interfaces and the vacuum boundary are corrected, trivially at this initial time, by the evolved density estimate in REMIX.

In Fig. H.2, we show the evolution of a simulation of a Jupiter-like planet with REMIX, but without the kernel normalising term. These demonstrate how the instabilities seen in Fig. 2.19 continue falling inwards at later times. As seen in Fig. 2.20, the densities of particles without the normalisation term are not tied to the distribution of mass in the simulation volume, leading to error driven instabilities. These grow to fully disturb the profile of the hydrogen–helium envelope, but are fully avoided in the full REMIX scheme.