Monthly Notices of the ROYAL ASTRONOMICAL SOCIETY

MNRAS **511,** 2367–2389 (2022) https://doi.org/10.1093/mnras/stab3166
Advance Access publication 2021 November 8

SPHENIX: smoothed particle hydrodynamics for the next generation of galaxy formation simulations

Josh Borrow [1,2★] Matthieu Schaller [3,4] Richard G. Bower [1] and Joop Schaye [4]

[1] *Institute for Computational Cosmology, Department of Physics, University of Durham, South Road, Durham DH1 3LE, UK*
[2] *Department of Physics, Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, Cambridge, MA 02139, USA*
[3] *Lorentz Institute for Theoretical Physics, Leiden University, PO Box 9506, NL-2300 RA Leiden, the Netherlands*
[4] *Leiden Observatory, Leiden University, PO Box 9513, NL-2300 RA Leiden, the Netherlands*

Accepted 2021 October 22. Received 2021 October 20; in original form 2020 December 7

ABSTRACT

Smoothed particle hydrodynamics (SPH) is a ubiquitous numerical method for solving the fluid equations, and is prized for its conservation properties, natural adaptivity, and simplicity. We introduce the SPHENIX SPH scheme, which was designed with three key goals in mind: to work well with sub-grid physics modules that inject energy, be highly computationally efficient (both in terms of compute and memory), and to be Lagrangian. SPHENIX uses a Density-Energy equation of motion, along with a variable artificial viscosity and conduction, including limiters designed to work with common sub-grid models of galaxy formation. In particular, we present and test a novel limiter that prevents conduction across shocks, preventing spurious radiative losses in feedback events. SPHENIX is shown to solve many difficult test problems for traditional SPH, including fluid mixing and vorticity conservation, and it is shown to produce convergent behaviour in all tests where this is appropriate. Crucially, we use the same parameters within SPHENIX for the various switches throughout, to demonstrate the performance of the scheme as it would be used in production simulations. SPHENIX is the new default scheme in the SWIFT cosmological simulation code and is available open source.

Key words: hydrodynamics – methods: numerical – galaxies: formation – galaxies: evolution.

1 INTRODUCTION

There have been many approaches to solving the equations of motion for a collisional fluid in a cosmological context over the years, from simple first-order fixed-grid (Cen 1992) to high-order discontinuous Galerkin schemes solved in an adaptive environment (Guillet et al. 2019). Because smoothed particle hydrodynamics (SPH) strikes the sweet spot between computational cost, stability, and adaptivity, it has been used throughout the astronomical community for nearly five decades.

SPH was originally developed by Lucy (1977) and Gingold & Monaghan (1977), and was used initially to model individual stars, as this problem was well suited to Lagrangian schemes. Shortly after, further applications of the method were deployed for the study of the fragmentation of gas clouds (Wood 1981), and for the formation of planets (Benz 1988).

The practical use of SPH in a cosmological context began with Hernquist & Katz (1989), which provided a novel solution to the large dynamic range of time-steps required to evolve a cosmological fluid, and was cemented by the GADGET-2 code (Springel 2005) that was made public and exploited worldwide to model galaxy formation processes within this context for the first time (e.g. Dolag et al. 2004; Ettori et al. 2006; Crain et al. 2007). The base SPH model released in GADGET-2, however, was relatively simple, consisting of a fixed artificial viscosity coefficient and scheme based on Monaghan (1992). Improved models existed, such as those presented in Monaghan (1997), but the key that led to the community rallying around GADGET-2 was both its open-source nature and scalability, with GADGET-2 able to run on hundreds or thousands of cores.

The popularity of GADGET-2, and similar codes like GASOLINE (Wadsley, Stadel & Quinn 2004), along with its relatively simple hydrodynamics model, led to critical works such as Agertz et al. (2007) and Bauer & Springel (2012) that pointed out flaws in their SPH modelling, relative to mesh-based codes of the time. The SPH community as a whole, however, already had solutions to these problems (see e.g. Price 2008) and many robust solutions were proposed and integrated into cosmological modelling codes. In Heß & Springel (2010), the authors experimented with an extension to GADGET-2 using a Voronoi mesh to reduce errors inherent in SPH and allow for better results on fluid mixing problems, eventually giving rise to the AREPO moving mesh scheme, allowing for significantly improved accuracy per particle but drastically increasing computational cost (Springel 2010; Weinberger, Springel & Pakmor 2020). In this case, the authors have steadily increased their computational cost per particle in an attempt to reduce errors inherent in their hydrodynamics model as much as practicable.

Other authors took different directions, with the GASOLINE code (Wadsley et al. 2004; Wadsley, Veeravalli & Couchman 2008; Wadsley, Keller & Quinn 2017) choosing to explicitly average pressures within the SPH equation of motion to alleviate the problems of artificial surface tension; the PHANTOM developers (Price 2008,

★ E-mail: josh@joshborrow.com

2012; Price et al. 2018) advocating for artificial conduction of energy; and further developments on the GADGET-2 and updated GADGET-3 code by Hopkins (2013) and Hu et al. (2014) based on the work by Saitoh & Makino (2013), using an explicit smoothed pressure scheme to ensure a consistent pressure field over the contact discontinuities that artificial surface tension arises from.

Simultaneously, there was a work to reduce the fundamental numerical errors present in SPH taking place by (Cullen & Dehnen 2010; Read, Hayfield & Agertz 2010; Dehnen & Aly 2012; Read & Hayfield 2012) through the use of improved choices for the SPH kernel, which up until this point was assumed to have a little effect on results from SPH simulations. These improved kernels typically have larger 'wings', encompassing more neighbours and providing more accurate reconstructions for smoothed quantities. These more accurate reconstructions are particularly important for the construction of accurate gradients, which enter into 'switches' that control the strength of the artificial viscosity and conduction terms.

The rise of more complex SPH models occurred alongside a significant jump in the complexity of the corresponding galaxy formation models; such an increase in complexity was required as resolutions increased over time, meaning more physics could be modelled directly. Many astrophysical processes take place on scales smaller than what can be resolved in simulations and are included in these so-called galaxy formation 'sub-grid' models. These processes include radiative cooling, which has progressed from a simple one parameter model to element and even ionization state-dependent rates (see e.g. Wiersma et al. 2009; Ploeckinger & Schaye 2020); star formation (see e.g. Cen & Ostriker 1992; Schaye & Dalla Vecchia 2008, and references therein); and stellar feedback to model supernovae and element outflows (see e.g. Navarro & White 1993; Springel & Hernquist 2003; Dalla Vecchia & Schaye 2008, 2012, and references therein). The coupling of these processes to hydrodynamics is complex and often overlooked; careful treatment of conservation laws and quirks of the chosen variables used to represent the fluid can frequently hide errors in plain sight (Borrow, Schaller & Bower 2020).

The development of the SWIFT code (Schaller et al. 2016) led to a re-implementation of the sub-grid model used for the EAGLE simulation (Schaye et al. 2015), and a chance to re-consider the ANARCHY SPH scheme that was used in the original (GADGET-3 based) code (see Schaller et al. 2015, for details on the scheme). The findings in Oppenheimer et al. (2018; their appendix D) and Borrow et al. (2020) meant that a switch away from the original Pressure-Entropy scheme to one based on a smoothed density field was preferred, along with the key design goals outlined below. This work describes the SPHENIX[1] scheme and demonstrates its performance on many hydrodynamics tests. We note here that SPHENIX does not give the best performance-per-particle (in both absolute values of L1 norm (see Section 5.1 for our definition of the L1 norm) compared to the analytical solution, and in terms of convergence speed) compared to other schemes. The moving mesh AREPO (Springel 2010), finite-volume GIZMO (Hopkins 2015), and corrected scheme presented in Rosswog (2020a) will produce improved results. SPHENIX however lies in the very low-cost (memory and computation) per-particle sweet-spot that traditional SPH schemes occupy, whilst maximizing performance with some novel limiters for artificial conduction and viscosity. This makes it an excellent choice for the default SPH scheme in SWIFT.

[1]Note that, similar to the popular GIZMO schemes, SPHENIX is not an acronym.

The remainder of this paper is organized as follows: in Section 2, we describe the SWIFT cosmological simulation code and the time-stepping algorithms present within it. In Section 3, we describe SPHENIX in its entirety. In Section 4, we describe the artificial conduction limiter used for energetic feedback schemes. Finally, in Section 5, we show how SPHENIX performs on various hydrodynamics problems.

2 THE SWIFT SIMULATION CODE

The SWIFT[2] simulation code (Schaller et al. 2016, 2018) is a hybrid parallel SPH and gravity code, designed to run across multiple compute nodes using the Message Passing Interface (MPI), but to utilize threads on each node (rather than the traditional method of using one MPI rank per core). This, along with its task-based parallelism approach, asynchronous communication scheme, and work-splitting domain decomposition system allow for excellent strong- and weak-scaling characteristics (Borrow et al. 2018).

SWIFT is also designed to be hugely modular, with hydrodynamics schemes, gravity schemes, and sub-grid models able to be easily swapped out. SWIFT can be configured to use a replica of the GADGET-2 hydrodynamics scheme (Springel & Hernquist 2002), a simplified version of the base PHANTOM scheme (Price et al. 2018), the MFM and MFV schemes described in Hopkins (2015), SPHENIX, or a host of other schemes. It can also be configured to use multiple different galaxy formation sub-grid models, including a very basic scheme (constant Λ cooling, no star formation), the EAGLE sub-grid model (Schaye et al. 2015), a 'Quick Lyman-α' model, the GEAR sub-grid model (Revaz & Jablonka 2012), and some further evolutions including cooling tables from Ploeckinger & Schaye (2020). The gravity solver is interchangeable but the one used here, and throughout all SWIFT simulations, uses the Fast Multipole Method (Greengard & Rokhlin 1987) with an adaptive opening angle, similar to Dehnen (2014).

2.1 Time integration

SWIFT uses a Velocity-Verlet scheme to integrate particles through time. This takes their acceleration ($\vec{a}$) from the equation of motion and time-step (Δt) and integrates their position forward in time through a Kick-Drift-Kick scheme as follows:

$$\vec{v}\left(t+\frac{\Delta t}{2}\right) = \vec{v}(t) + \tfrac{\Delta t}{2}\vec{a}(t), \tag{1}$$

$$\vec{r}\,(t+\Delta t) = \vec{r}(t) + \vec{v}\left(t+\tfrac{\Delta t}{2}\right)\Delta t, \tag{2}$$

$$\vec{v}\,(t+\Delta t) = \vec{v}\left(t+\tfrac{\Delta t}{2}\right) + \tfrac{\Delta t}{2}\vec{a}(t+\Delta t), \tag{3}$$

where the first and last equations, updating the velocity, are referred to as the 'kick', and the central equation is known as the 'drift'. The careful observer will note that the 'drift' can be split into as many pieces as required allowing for accurate interpolation of the particle position in-between kick steps. This is important in cosmological galaxy formation simulations, where the dynamic range is large. In

[2]For the interested reader, the implementation of the SPHENIX scheme was developed fully in the open and is available in the SWIFT repository at http://swiftsim.com (Schaller et al. 2018), including all of the tests and examples shown below. We use version 0.9.0 of the SWIFT code for the tests in this work.

this case, particles are evolved with their own, particle-carried time-step, given by

$$\Delta t_i = \mathrm{C_{CFL}} \frac{2\gamma_K h_i}{v_{\mathrm{sig},i}}, \tag{4}$$

dependent on the Courant–Friedrichs–Lewy ($\mathrm{C_{CFL}}$; Courant, Friedrichs & Lewy 1928) constant, the kernel-dependent relationship between cut-off and smoothing length γ_K, particle-carried smoothing length h_i, and signal velocity $v_{\mathrm{sig},i}$ (see equation 30). The discussion of the full time-stepping algorithm is out of the scope of this work, but see Hernquist & Katz (1989) and Borrow, Vandenbroucke & Schaller (2019) for more information.

2.1.1 Time-step limiter

As the time-step of the particles is particle-carried, there may be certain parts of the domain that contain interacting particles with vastly different time-steps (this is particularly promoted by particles with varied temperatures within a given kernel). Having these particles interact is problematic for a number of reasons, and as such we include the time-step limiter described in Saitoh & Makino (2009) Durier & Dalla Vecchia (2012) in all problems solved below. SWIFT chooses to limit neighbouring particles to have a maximal time-step difference of a factor of 4.

3 SPHENIX

The SPHENIX scheme was designed to replace the ANARCHY scheme used in the original EAGLE simulations for use in the SWIFT simulation code. This scheme had three major design goals:

(i) Be a Lagrangian SPH scheme, as this has many advantages and is compatible with the EAGLE sub-grid model.
(ii) Work well with the EAGLE sub-grid physics, namely instantaneous energy injection and sub-grid cooling.
(iii) Be highly computationally and memory efficient.

The last requirement precludes the use of any Riemann solvers in so-called GIZMO-like schemes (although these do not necessarily give improved results for astrophysical problem sets, see Borrow et al. 2019); see Appendix A. The second requirement also means that the use of a pressure-based scheme (such as ANARCHY) is not optimal, see Borrow et al. (2020) for more details.

The SPHENIX scheme is based on so-called 'Traditional' Density-Energy SPH. This means that it uses the smoothed mass density,

$$\hat{\rho}(\vec{x}) = \sum_j m_j W(|\vec{x} - \vec{x}_j|, h(\vec{x})) \tag{5}$$

where here j are indices describing particles in the system, $h(\vec{x})$ is the smoothing length evaluated at position $\vec{x}$, and $W(r, h)$ is the kernel function.

In the examples below, the Quartic Spline (M5) kernel,

$$w(q) = \begin{cases} \left(\frac{5}{2} - q\right)^4 - 5\left(\frac{3}{2} - q\right)^4 + 10\left(\frac{1}{2} - q\right)^4 & q < \frac{1}{2} \\ \left(\frac{5}{2} - q\right)^4 - 5\left(\frac{3}{2} - q\right)^4 & \frac{1}{2} \leq q < \frac{3}{2} \\ \left(\frac{5}{2} - q\right)^4 & \frac{3}{2} \leq q < \frac{5}{2} \\ 0 & q \geq \frac{5}{2} \end{cases} \tag{6}$$

with $W(r, h) = \kappa_{n_D} w(r/h)/h^{n_D}$, n_D the number of dimensions, and $\kappa_3 = (7/478\pi)$ for three dimensions, is used. The SPHENIX scheme has also been tested with other kernels, notably the Cubic and Quintic Spline (M4, M6) and the Wendland (C2, C4, C6) kernels (Wendland 1995). The choice of kernel does not qualitatively affect the results in any of the tests in this work (see Dehnen & Aly 2012, for significantly more information on kernels). Higher order kernels do allow for lower errors on tests that rely on extremely accurate reconstructions to cancel forces (for instance the Gresho–Chan vortex, Section 5.3), but we find that the Quintic Spline provides an excellent trade-off between computational cost and accuracy in practical situations. Additionally, the Wendland kernels do have the benefit that they are not susceptible to the pairing instability, but they must have an ad-hoc correction applied in practical use (Dehnen & Aly 2012, section 2.5). We find no occurrences of the pairing instability in both the tests and our realistic simulations. The SPHENIX scheme is kernel-invariant, and as such can be used with any reasonable SPH kernel.

The smoothing length h is determined by satisfying

$$\hat{n}(\vec{x}) = \sum_j W\left(|\vec{x} - \vec{x}_j|, h(\vec{x})\right) = \left(\frac{\eta}{h(\vec{x})}\right)^{n_D}, \tag{7}$$

with η setting the resolution scale. The precise choice for η generally does not qualitatively change results; here, we choose $\eta = 1.2$ due to this value allowing for a very low E_0 error (see Read et al. 2010; Dehnen & Aly 2012),[3] which is a force error originating from particle disorder within a single kernel. In SWIFT, these equations are solved numerically to a relative accuracy of 10^{-4}.

The smoothed mass density, along with a particle-carried internal energy per unit mass u, is used to determine the pressure at a particle position through the equation of state

$$P(\vec{x}_i) = P_i = (\gamma - 1) u_i \hat{\rho}_i, \tag{8}$$

with γ the ratio of specific heats, taken to be 5/3 throughout unless specified. This pressure enters the first law of thermodynamics,

$$\left.\frac{\partial u_i}{\partial \vec{q}_i}\right|_{A_i} = -\frac{P_i}{m_i}\frac{\partial V_i}{\partial \vec{q}_i}, \tag{9}$$

with $\vec{q}_i$ a state vector containing both $\vec{x}_i$ and h_i as *independent* variables, A_i the entropy of particle i (i.e. this equation only applies to dissipationless dynamics), and $V_i = m_i/\hat{\rho}_i$ describing the volume represented by particle i. This constraint, along with the one on the smoothing length, allows for an equation of motion to be extracted from a Lagrangian (see e.g. the derivations in Springel & Hernquist 2002; Hopkins 2013),

$$\frac{\mathrm{d}\vec{v}_i}{\mathrm{d}t} = -\sum_j m_j \left[\frac{f_{ij} P_i}{\hat{\rho}_i^2}\nabla_i W_{ij} + \frac{f_{ji} P_j}{\hat{\rho}_j^2}\nabla_j W_{ji}\right], \tag{10}$$

where $W_{ab} = W(|\vec{x}_b - \vec{x}_a|, h(\vec{x}_a))$, $\nabla_a = \partial/\partial\vec{x}_a$, and f_{ab} a dimensionless factor encapsulating the non-uniformity of the smoothing length field

$$f_{ab} = 1 - \frac{1}{m_b}\left(\frac{h_a}{n_D \hat{n}_a}\frac{\partial \hat{\rho}_i}{\partial h_i}\right)\left(1 + \frac{h_i}{n_D \hat{n}_i}\frac{\partial \hat{n}_i}{\partial h_i}\right)^{-1} \tag{11}$$

[3] This corresponds to ∼58 weighted neighbours for our Quartic Spline in a scenario where all neighbours have uniform smoothing lengths. In practical simulations the 'number of neighbours' that a given particle interacts with can vary by even orders of magnitude but equation (7) must be satisfied for all particles ensuring an accurate reconstruction of the field. More discussion on this choice of smoothing length can be found in (Monaghan 2002; Springel & Hernquist 2002; Price 2007, 2012; Borrow et al. 2020). We chose $\eta = 1.2$ based on fig. 3 in Dehnen & Aly (2012), where this corresponds to a very low reconstruction error in the density.

and is generally of order unity. There is also an associated equation of motion for internal energy,

$$\frac{\mathrm{d}u_i}{\mathrm{d}t} = -\sum_j m_j f_{ij} \frac{P_i}{\hat{\rho}_i^2} \vec{v}_{ij} \cdot \nabla_i W_{ij}, \tag{12}$$

with $\vec{v}_{ij} = \vec{v}_j - \vec{v}_i$. Note that other differences between vector quantities are defined in a similar way, including for the separation of two particles $\vec{x}_{ij} = \vec{x}_j - \vec{x}_i$.

3.1 Artificial viscosity

These equations, due to the constraint of constant entropy introduced in the beginning, lead to naturally dissipationless solutions; they cannot capture shocks. Shock capturing in SPH is generally performed using 'artificial viscosity'.

The artificial viscosity implemented in SPHENIX is a simplified and modified extension to the Cullen & Dehnen (2010) 'inviscid SPH' scheme. This adds the following terms to the equation of motion (see Monaghan 1992, and references within),

$$\frac{\mathrm{d}\vec{v}_i}{\mathrm{d}t} = -\sum_j m_j \zeta_{ij} \left[f_{ij} \nabla_i W_{ij} + f_{ji} \nabla_j W_{ji} \right], \tag{13}$$

and to the associated equation of motion for the internal energy,

$$\frac{\mathrm{d}u_i}{\mathrm{d}t} = -\frac{1}{2} \sum_j m_j \zeta_{ij} \vec{v}_{ij} \cdot \left[f_{ij} \nabla_i W_{ij} + f_{ji} \nabla_j W_{ji} \right], \tag{14}$$

where ζ_{ij} controls the strength of the viscous interaction. Note here that the internal energy equation of motion is explicitly symmetrized, which was not the case for the SPH equation of motion for internal energy (equation 12). In this case, which means that there are terms from both the *ij* and *ji* interactions in equation (14), whereas in equation (12), there is only a term from the *ij* interaction. This choice was due to the symmetric version of this equation performing significantly better in the test examples below, likely due to multiple time-stepping errors within regions where the viscous interaction is the strongest.[4]

There are many choices available for ζ_{ij}, with the case used here being

$$\zeta_{ij} = -\alpha_V \mu_{ij} \frac{v_{\mathrm{sig},ij}}{\hat{\rho}_i + \hat{\rho}_j}, \tag{15}$$

where

$$\mu_{ij} = \begin{cases} \frac{\vec{v}_{ij} \cdot \vec{x}_{ij}}{|\vec{x}_{ij}|} & \vec{v}_{ij} \cdot \vec{x}_{ij} < 0 \\ 0 & \vec{v}_{ij} \cdot \vec{x}_{ij} \geq 0 \end{cases} \tag{16}$$

is a basic particle-by-particle converging flow limiter (meaning that the viscosity term vanishes when $\nabla \cdot \vec{v} \geq 0$), and

$$v_{\mathrm{sig},ij} = c_i + c_j - \beta_V \mu_{ij}, \tag{17}$$

is the signal velocity between particles i and j, with $\beta_V = 3$, a dimensionless constant, and with c_i the soundspeed of particle i defined through the equation of state as

$$c_i = \sqrt{\frac{P_i}{\hat{\rho}_i}} = \sqrt{(\gamma - 1)\gamma u_i}. \tag{18}$$

Finally, the dimensionless viscosity coefficient α_V (Monaghan & Gingold 1983) is frequently taken to be a constant of order unity. In SPHENIX, this becomes an interaction-dependent constant (see Morris & Monaghan 1997; Cullen & Dehnen 2010, for similar schemes), with $\alpha_V = \alpha_{V,ij}$, dependent on two particle-carried α values as follows:

$$\alpha_{V,ij} = \frac{1}{4}(\alpha_{V,i} + \alpha_{V,j})(B_i + B_j), \tag{19}$$

where

$$B_i = \frac{|\nabla \cdot \vec{v}_i|}{|\nabla \cdot \vec{v}_i| + |\nabla \times \vec{v}_i| + 10^{-4} c_i / h_i} \tag{20}$$

is the Balsara (1989) switch for particle i, which allows for the deactivation of viscosity in shear flows, where there is a high value of $\nabla \cdot \vec{v}$, but the associated shear viscosity is unnecessary. This, in particular, affects rotating shear flows, such as galaxy discs, where the scheme used to determine $\alpha_{V,i}$ described below will return a value close to the maximum.

The equation for $\alpha_{V,i}$ is solved independently for each particle over the course of the simulation. Note that $\alpha_{V,i}$ is never drifted, and is only ever updated at the 'kick' steps. The source term in the equation for $\alpha_{V,i}$, designed to activate the artificial viscosity within shocking regions, is the shock indicator

$$S_i = \begin{cases} -h_i^2 \max\left(\dot{\nabla} \cdot \vec{v}_i, 0\right) & \nabla \cdot \vec{v}_i \leq 0 \\ 0 & \nabla \cdot \vec{v}_i > 0 \end{cases} \tag{21}$$

where, here the time differential of the local velocity divergence field

$$\dot{\nabla} \cdot \vec{v}_i(t + \Delta t) = \frac{\nabla \cdot \vec{v}_i(t + \Delta t) - \nabla \cdot \vec{v}_i(t)}{\Delta t} \tag{22}$$

with $\nabla \cdot \vec{v}_i$ the local velocity divergence field and Δt the time-step associated with particle i. The primary variable in the shock indicator S_i of $\dot{\nabla} \cdot \vec{v}$ is high in pre-shock regions, with the secondary condition for the flow being converging ($\nabla \cdot \vec{v} \leq 0$) helpful to avoid false detections as the Balsara (1989) switch is used independently from the equation that evolves $\alpha_{V,i}$ (this choice is notably different from most other schemes that use B_i directly in the shock indicator S_i). This choice allows for improved shock capturing in shearing flows (e.g. feedback events occurring within a galaxy disc). In these cases, the Balsara (1989) switch (which is instantaneously evaluated) rapidly becomes close to 1.0, and the already high value of $\alpha_{V,i}$ allows for a strong viscous reaction from the fluid. The shock indicator is then transformed into an optimal value for the viscosity coefficient as

$$\alpha_{V,\mathrm{loc},i} = \alpha_{V,\max} \frac{S_i}{c_i^2 + S_i}, \tag{23}$$

with a maximum value of $\alpha_{V,\max} = 2.0$ for $\alpha_{V,\mathrm{loc}}$. The value of $\alpha_{V,i}$ is then updated as follows:

$$\alpha_{V,i} = \begin{cases} \alpha_{V,\mathrm{loc},i} & \alpha_{V,i} < \alpha_{V,\mathrm{loc},i} \\ \frac{\alpha_{V,i} + \alpha_{V,\mathrm{loc},i} \frac{\Delta t}{\tau_{V,i}}}{1 + \frac{\Delta t}{\tau_{V,i}}} & \alpha_{V,i} > \alpha_{V,\mathrm{loc},i} \end{cases} \tag{24}$$

where $\tau_{V,i} = \gamma_K \ell_V h_i / c_i$ with γ_K the 'kernel gamma', a kernel dependent quantity relating the smoothing length and compact support ($\gamma_K = 2.018932$ for the quartic spline in 3D; Dehnen & Aly 2012) and ℓ_V, a constant taking a value of 0.05. The final value of $\alpha_{V,i}$ is checked against a minimum, however, the default value of this minimum is zero and the evolution strategy used above guarantees that $\alpha_{V,i}$ is strictly positive and that the decay is stable regardless of time-step.

3.2 Artificial conduction

Attempting to resolve sharp discontinuities in non-smoothed variables in SPH leads to errors. This can be seen above with strong

[4]For these reasons all recent works choose symmetric forms for these equations.

velocity discontinuities (shocks) not being correctly handled and requiring an extra term in the equation of motion (artificial viscosity) to be captured. A similar issue arises when attempting to resolve strong discontinuities in internal energy (temperature). To resolve this, we introduce an artificial energy conduction scheme similar to the one presented by Price (2008). This adds an extra term to the equation of motion for internal energy,

$$\frac{\mathrm{d}u_i}{\mathrm{d}t} = \sum_j m_j v_{D,ij}(u_i - u_j)\hat{r}_{ij} \cdot \left(f_{ij}\frac{\nabla_i W_{ij}}{\hat{\rho}_i} + f_{ji}\frac{\nabla_j W_{ji}}{\hat{\rho}_j} \right) \quad (25)$$

with $\hat{r}_{ij}$ the unit vector between particles i and j, and where

$$v_{D,ij} = \frac{\alpha_{D,ij}}{2}\left(\frac{|\vec{v}_{ij}\cdot\vec{x}_{ij}|}{|\vec{x}_{ij}|} + \sqrt{2\frac{|P_i - P_j|}{\hat{\rho}_j + \hat{\rho}_j}} \right). \quad (26)$$

This conductivity speed is the average of two commonly used speeds, with the former velocity-dependent term taken from Price et al. (2018; modified from Wadsley et al. 2008), and the latter pressure-dependent term taken from Price (2008). These are usually used separately for cases that aim to reduce entropy generation in disordered fields and contact discontinuities, respectively (where initially there is a strong discontinuity in pressure that is removed by the artificial conduction scheme), but we combine them here as both cases are relevant in galaxy formation simulations and use this same velocity throughout our testing, a notable difference with other works using conduction schemes (e.g. Price et al. 2018). Price et al. (2018) avoided pressure-based terms in simulations with self-gravity, but they use no additional terms (e.g. our α_D) to limit conduction in flows where it is not required. This is additionally somewhat similar to the conduction speed used in ANARCHY and Hu et al. (2014), which is a modified version of the signal velocity (equation 17) with our speed replacing the sum of sound speeds with a differenced term. Appendix B contains an investigation of the individual terms in the conduction velocity. The interaction-dependent conductivity coefficient,

$$\alpha_{D,ij} = \frac{P_i\alpha_{D,i} + P_j\alpha_{D,j}}{P_i + P_j}, \quad (27)$$

is pressure-weighted to enable the higher pressure particle to lead the conduction interaction, a notable departure from other thermal conduction schemes in use today. This is critical when it comes to correctly applying the conduction limiter during feedback events described below. The individual particle-carried $\alpha_{D,i}$ are ensured to only be active in cases where there is a strong discontinuity in internal energy. This is determined by using the following discontinuity indicator,

$$K_i = \beta_D \gamma_K h_i \frac{\nabla^2 u_i}{\sqrt{u_i}}, \quad (28)$$

where β_D is a fixed dimensionless constant taking a value of 1. The discontinuity indicator enters the time differential for the individual conduction coefficients as a source term,

$$\frac{\mathrm{d}\alpha_{D,i}}{\mathrm{d}t} = K_i + \frac{\alpha_{D,\min} - \alpha_{D,i}}{\tau_{D,i}}, \quad (29)$$

with $\tau_{D,i} = \gamma_K h_i / v_{\mathrm{sig},i}$, $\alpha_{D,\min} = 0$ the minimal allowed value of the artificial conduction coefficient, and with the individual particle signal velocity,

$$v_{\mathrm{sig},i} = \max_j(v_{\mathrm{sig},ij}), \quad (30)$$

controlling the decay rate. $\nabla^2 u$ is used as the indicator for a discontinuity, as opposed to ∇u, as it allows for (physical, well represented within SPH) linear gradients in internal energy to be maintained without activating artificial conduction. This is then integrated during 'kick' steps using

$$\alpha_{D,i}(t + \Delta t) = \alpha_{D,i}(t) + \frac{\mathrm{d}\alpha_{D,i}}{\mathrm{d}t}\Delta t. \quad (31)$$

The final stage of evolution for the individual conduction coefficients is to limit them based on the local viscosity of the fluid. This is necessary because thermal feedback events explicitly create extreme discontinuities within the internal energy field that lead to shocks (see Section 4 for the motivation leading to this choice). The limit is applied using the maximal value of viscous alpha among the neighbours of a given particle,

$$\alpha_{V,\max,i} = \max_j(\alpha_{V,j}), \quad (32)$$

with the limiter being applied using the maximally allowed value of the conduction coefficient,

$$\alpha_{D,\max,i} = \alpha_{D,\max}\left(1 - \frac{\alpha_{V,\max,i}}{\alpha_{V,\max}}\right), \quad (33)$$

with $\alpha_{D,\max} = 1$ a constant, and

$$\alpha_{D,i} = \begin{cases} \alpha_{D,i} & \alpha_{D,i} < \alpha_{D,\max} \\ \alpha_{D,\max} & \alpha_{D,i} > \alpha_{D,\max}. \end{cases} \quad (34)$$

This limiter allows for a more rapid increase in conduction coefficient, and a higher maximum, than would usually be viable in simulations with strong thermal feedback implementations. In ANARCHY, another scheme employing artificial conduction, the rate at which the artificial conduction could grow was chosen to be significantly smaller. In ANARCHY, $\beta_D = 0.01$, which is 100 times smaller than the value chosen here (Schaye et al. 2015, appendix A3). This additional conduction is required to accurately capture contact discontinuities with a Density-Energy SPH equation of motion.

4 MOTIVATION FOR THE CONDUCTION LIMITER

The conduction limiter first described in Section 3 is formed of two components; a maximal value for the conduction coefficient in viscous flows (equation 34), and one that ensures that a particle with a higher pressure takes preference within the conduction interaction (equation 27).

This limiter is necessary due to interactions of the artificial conduction scheme with the sub-grid physics model. Here, the EAGLE sub-grid model is shown as this is what SPHENIX was designed for use with, however all schemes employing energetic feedback and unresolved cooling times will suffer from the same problems when using an un-limited artificial conduction. In short, when an energetic feedback event takes place, the artificial conduction switch is activated (as this is performed by injecting lots of energy into one particle, leading to an extreme value of $\nabla^2 u$). This then leads to energy leaking out of the particle ahead of the shock front, which is then radiated away as the neighbouring particles can rapidly cool due to their temperature being lower leading to smaller cooling times.

To show the effect of this problem on a real system, we set up a uniform volume containing 32^3 gas particles at approximately solar metallicity ($Z = 0.014$) and equilibrium temperature (around 10^4 K), at a density of $n_{\mathrm{H}} = 0.1$ cm^{-3}. The central particle in the volume has approximately the same amount of energy injected into it as in a single EAGLE-like stellar feedback event (heating it to $\sim 10^{7.5}$ K) at the start of the simulation and the code is ran with full sub-grid cooling (using the tables from Wiersma et al. 2009) enabled. The

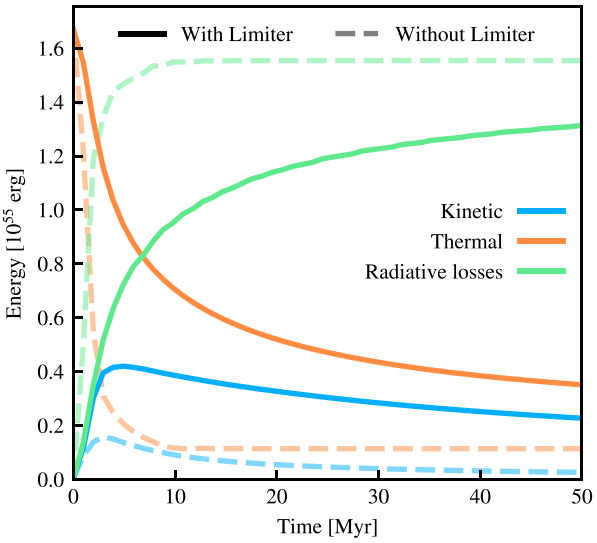

Figure 1. Energy in various components as a function of time for a simulated supernova blast (see text for details of the set-up). Blue shows energy in the kinetic phase, orange shows energy in the thermal phase (neglecting the thermal energy of the background), and green shows energy lost to radiation. The solid lines show the simulation performed with the artificial conduction limiter applied and the dashed lines show the simulation without any such limiter. Simulations performed without the limiter show huge, rapid, and cooling losses.

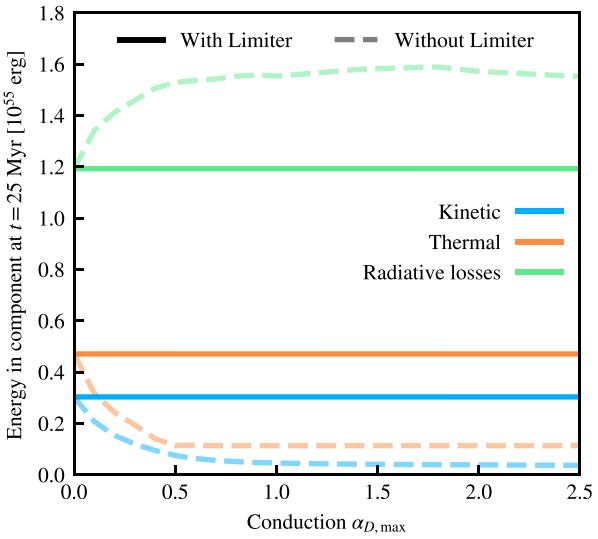

Figure 2. The set-up from Fig. 1 performed for different values for the maximum artificial conduction coefficient $\alpha_{D,\max}$ (i.e. a different horizontal axis as Fig. 1, with the same vertical axis), now showing the components of energy in each phase at a fixed time of $t = 25$ Myr. Colours and line styles are the same as in Fig. 1. As well as demonstrating the issue with an un-limited conduction, this figure shows that the conduction limiter prevents the loss of additional energy relative to a simulation performed without any artificial conduction.

initial values for the artificial viscosity and conduction coefficients are set to be zero (whereas in practice they are set to be their maximum and minimum in 'real' feedback events; this has a little effect on the results as the coefficients rapidly stabilize).

Fig. 1 shows the energy in the system (with the thermal energy of the 'background' particles removed to ensure a large dynamic range in thermal energy is visible on this plot) in various components. We see that, at least for the case with the limiter applied, at $t = 0$, there is the expected large injection of thermal energy that is rapidly partially transformed into kinetic energy as in a classic blastwave problem (like the one shown in Fig. 5; in our idealized, non-radiative, Sedov blasts only 28 per cent of the injected thermal energy is converted to kinetic energy). A significant fraction, around two thirds, of the energy is lost to radiation, but the key here is that there is a transformation of the initial thermal injection to a kinetic wave.

In the same simulation, now with the conduction limiter removed (dashed lines), almost all of the injected energy is immediately lost to radiation (i.e. the feedback is unexpectedly inefficient). The internal energy in the affected particle is rapidly conducted to its neighbours (that are then above, but closer to, the equilibrium temperature) which have a short cooling time and hence the energy is quickly lost.

The direct effect of the conduction limiter is shown in Fig. 2, where the same problem as above is repeated 10 times with maximal artificial conduction coefficients $\alpha_{D,\max}$ of 0–2.5 in steps of 0.1 (note that the value of $\alpha_{D,\max}$ used in production simulations is 1). We choose to show these extreme values to demonstrate the efficacy of the limiter even in extreme scenarios. The simulations with and without the limiter show the same result at $\alpha_{D,\max} = 0$ (i.e. with conduction disabled) but those without the limiter show a rapidly increasing fraction of the energy lost to cooling as the maximal conduction coefficient increases. The simulations with the limiter show a stable fraction of energy (at this fixed time of $t = 25$ Myr) in each component, showing that the limiter is working as expected and is curtailing these numerical radiative losses. This result is qualitatively unchanged for a factor of 100 higher, or lower, density background gas (i.e. gas between $n_{\rm H} = 0.001$ cm^{-3} and $n_{\rm H} = 10.0$ cm^{-3}). In both of these cases, the conduction can rapidly cause numerical radiative losses, but with the limiter enabled this is remedied entirely. We also note that the limiter remains effective even for extreme values of the conduction parameter (e.g. with $\alpha_{D,\max} = 100$), returning the same result as the case without artificial conduction for this test.

5 HYDRODYNAMICS TESTS

In this section, the performance of SPHENIX is shown on hydrodynamics tests, including the Sod (1978) shock tube, Sedov (1959) blastwave, and the Gresho & Sani (1990) vortex, along with many other problems relevant to galaxy formation. All problems are performed in hydrodynamics-only mode, with no radiative cooling or any other additional physics, and all use a $\gamma = 5/3$ equation of state ($P = (2/3)u_i\hat{\rho}$).

Crucially, all tests were performed with the same scheme parameters and settings, meaning that all of the switches are consistent (even between self-gravitating and pure hydrodynamical tests) unless otherwise stated. This departs from most studies where parameters are set for each problem independently, in an attempt to demonstrate the maximal performance of the scheme for a given test. The parameters used are as follows:

(i) The quartic spline kernel.

(ii) CFL condition $C_{\rm CFL} = 0.2$, with multiple time-stepping enabled (see e.g. Lattanzio et al. 1986).

(iii) Viscosity alpha $0.0 \leq \alpha_V \leq 2.0$ with the initial value being $\alpha_V = 0.1$ (similar to Cullen & Dehnen 2010).

(iv) Viscosity beta $\beta_V = 3.0$ and length $\ell_V = 0.05$ (similarly to Cullen & Dehnen 2010).

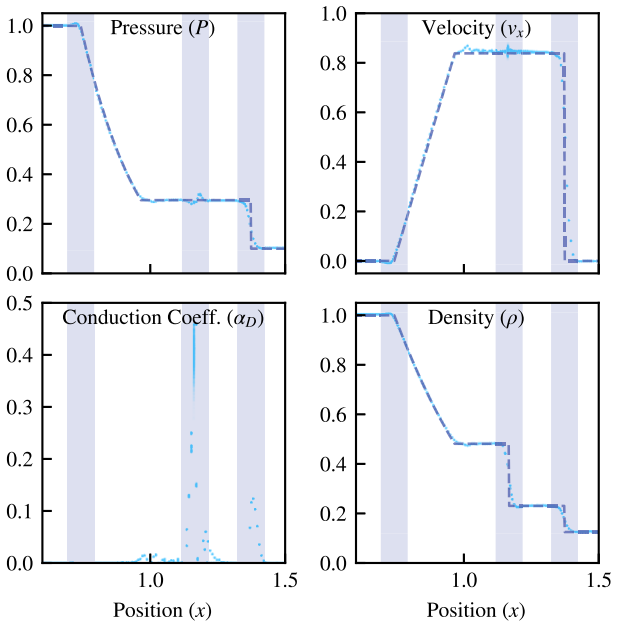

Figure 3. Individual quantities plotted against the analytic solution (purple dashed line) for the Sod shock tube in 3D. The horizontal axis shows the x position of the particles. All particles are shown in blue, with the purple shading in the background showing the regions considered for the convergence (Fig. 4) with the rarefaction wave, contact discontinuity, and shock, shown from left to right. All panels are shown at the same time $t = 0.2$, and for the same resolution level, using the 64^3 and 128^3 initial conditions for $x < 1$ and $x > 1$, respectively.

(v) Conduction alpha $0.0 \leq \alpha_D \leq 1.0$ (a choice consistent with Price 2008) with the viscosity-based conduction limiter enabled and the same functional form for the conduction speed (equation 26) used in all simulations.

(vi) Conduction beta $\beta_D = 1.0$ with the initial value of $\alpha_D = 0.0$.

These choices were all 'calibrated' to achieve an acceptable result on the Sod shock tube, and then left fixed with the results from the rest of the tests left unseen. We choose to present the tests in this manner in an effort to show a representative overview of the performance of SPHENIX in real-world conditions as it is primarily designed for practical use within future galaxy formation simulations.

The source code required to produce the initial conditions (or a link to download the initial conditions themselves if this is impractical) are available open source from the SWIFT repository.

5.1 Sod shock tube

The Sod (1978) shock tube is a classic Riemann problem often used to test hydrodynamics codes. The tube is made up of three main sections in the final solution: the rarefaction wave (between $0.7 < x < 1.0$), contact discontinuity (at position $x \approx 1.2$), and a weak shock (at position $x \approx 1.4$) at the time that we show it in Fig. 3.

5.1.1 Initial conditions

The initial conditions for the Sod shock tube uses body-centred cubic lattices to ensure maximally symmetric lateral forces in the initial state. Two lattices with equal particle masses, one at a higher density by a factor of 8 (e.g. one with 32^3 particles and one with 64^3 particles) are attached at $x = 1.0$.[5] This forms a discontinuity, with the higher density lattice being placed on the left with $\rho_L = 1$ and the lower density lattice on the right with $\rho_R = 1/8$. The velocities are initially set to zero for all particles and pressures set to be $P_L = 1$ and $P_R = 0.1$.

5.1.2 Results

Fig. 3 shows the shock tube at $t = 1$, plotted against the analytic solution. This figure shows the result from the 64^3 and 128^3 initial condition. In general, the simulation data (blue points) shows very close agreement with the analytic solution (purple dashed line).

The three purple bands correspond to three distinct regions within the shock tube. The furthest left is the rarefaction wave, which is an adiabatically expanding fluid. The band covers the turnover point of the wave, as this is where the largest deviation from the analytic solution is present. There is a slight overestimation of the density at this turnover point, primarily due to the symmetric nature of the SPH kernel.

The next band shows the contact discontinuity. No effort is made to suppress this discontinuity in the initial conditions (i.e. they are not relaxed). There is a small pressure blip of a similar size to that seen with schemes employing Riemann solvers, such as GIZMO (Hopkins 2015). There is no large velocity discontinuity associated with this pressure blip as is seen with SPH schemes that do not explicitly treat the contact discontinuity (note that every particle present in the simulation is shown here) with some form of conduction, a smoothed pressure field, or other method. Due to the strong discontinuity in internal energy present in this region, the artificial conduction coefficient α_D peaks, allowing for the pressure 'blip' to be reduced to one with a linear gradient.

The final section of the system, the rightmost region, is the shock. This shock is well captured by the scheme. There is a small activation of the conduction coefficient in this region, which is beneficial as it aids in stabilizing the shock front (Hu et al. 2014). This shows that the conduction limiter (Section 4) does not eliminate this beneficial property of artificial conduction within these frequently present weak (leading to $\alpha_V \lesssim 1.0$) shocks.

In an ideal case, the scheme would be able to converge at second order $L_1 \propto h^2$ away from shocks, and at first order $L_1 \propto h$ within shocks (Price et al. 2018). Here the L_1 norm of a band is defined as

$$L_1(K) = \frac{1}{n} \sum_n |K_{\rm sim}(\vec{x}) - K_{\rm ref}(\vec{x})| \tag{35}$$

with K some property of the system such as pressure, the subscripts sim and ref referring to the simulation data and reference solution, respectively, and n the number of particles in the system.

Fig. 4 shows the convergence properties of the SPHENIX scheme on this problem, using the pressure field in this case as the convergence variable. Compared to a scheme without artificial conduction (dotted lines), the SPHENIX scheme shows significantly improved convergence and a lower norm in the contact discontinuity, without sacrificing accuracy in other regions.

[5]This simplistic particle arrangement does cause a slight problem at the interface at higher (i.e. greater than one) dimensions. In 3D, some particles may have spurious velocities in the y and z directions at the interface, due to asymmetries in the neighbours found on the left and right side of the boundary. To offset this, the lattices are placed so that the particles are aligned along the x-axis wherever possible over the interface, however some spurious forces still result.

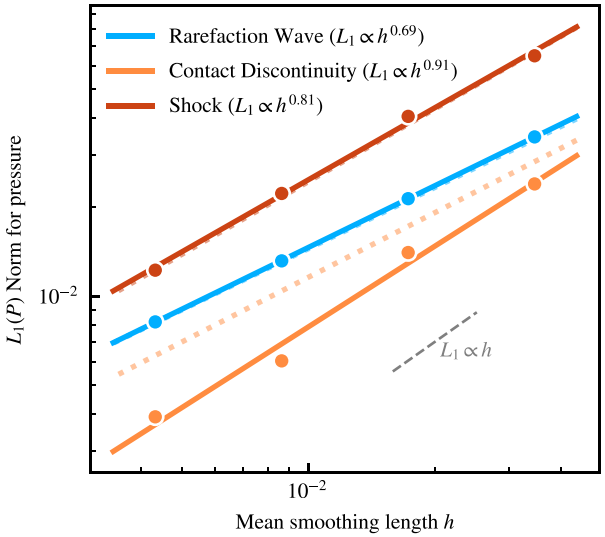

Figure 4. Pressure convergence for the three regions in Fig 3. The solid lines show fits to the data at various resolution levels (points) for each region, with the dotted lines showing convergence speed when the artificial conduction term is removed. The dashed grey line shows the expected speed of convergence for shocks in SPH simulations, to guide the eye, with a dependence of $L_1 \propto h$.

5.2 Sedov–Taylor blastwave

The Sedov–Taylor blastwave (Sedov blast; Taylor 1950; Sedov 1959) follows the evolution of a strong shock front through an initially isotropic medium. This is a highly relevant test for cosmological simulations, as this is similar to the implementations used for sub-grid (below the resolution scale) feedback from stars and black holes. In SPH schemes, this effectively tests the artificial viscosity scheme for energy conservation; if the scheme does not conserve energy, the shock front will be misplaced.

5.2.1 Initial conditions

Here, we use a glass file generated by allowing a uniform grid of particles to settle to a state, where the kinetic energy has stabilized. The particle properties are then initially set such that they represent a gas with adiabatic index $\gamma = 5/3$, a uniform pressure of $P_0 = 10^{-6}$, density $\rho_0 = 1$, all in a 3D box of side-length 1. Then, the $n = 15$ particles closest to the centre of the box have energy $E_0 = 1/n$ injected into them. This corresponds, roughly, to a temperature jump of a factor of $\sim 10^5$ over the background medium.

5.2.2 Results

Fig. 5 shows the particle properties of the highest resolution initial condition (128^3) at $t = 0.1$ against the analytic solution. The SPHENIX scheme closely matches the analytic solution in all particle fields, with the only deviation (aside from the smoothed shock front, an unavoidable consequence of using an SPH scheme) being a slight upturn in pressure in the central region (due to a small amount of conduction in this region). Of particular note is the position of the shock front matching exactly with the analytic solution, showing that the scheme conserves energy in this highly challenging situation thanks to the explicitly symmetric artificial viscosity equation of motion. The SPHENIX scheme shows qualitatively similar results to the PHANTOM scheme on this problem (Price et al. 2018).

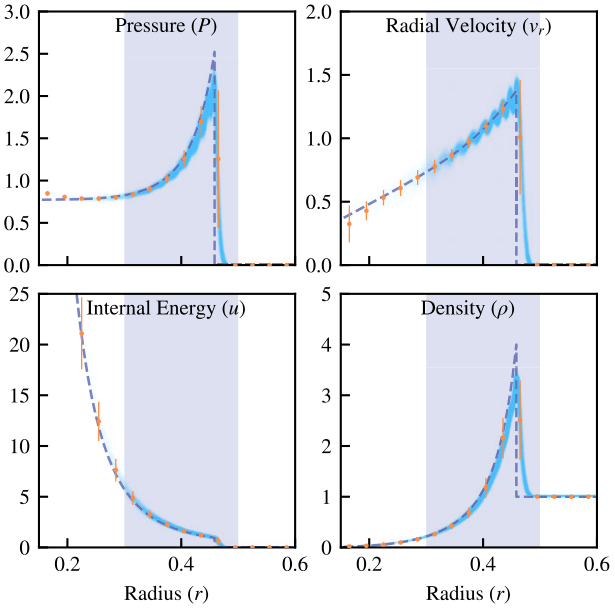

Figure 5. Particle properties at $t = 0.1$ shown against the analytic solution (purple dashed line) for the Sedov–Taylor blastwave. A random sub-set of 1/5th of the particles are shown in blue, with the orange points showing the mean value within equally spaced horizontal bins with one standard deviation of scatter. The background purple band shows the region considered for measuring convergence in Fig. 6. This figure shows the results for a 128^3 particle glass file.

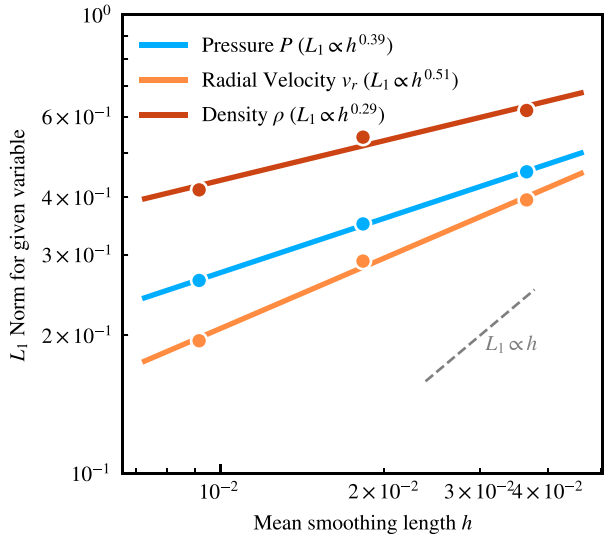

Figure 6. L_1 Convergence with mean smoothing length for various particle fields in the Sedov–Taylor blastwave test, measured at $t = 0.1$ against the analytic solution within the purple band of Fig. 5. Each set of points shows a measured value from an individual simulation, with the lines showing a linear fit to the data in logarithmic space. Dotted lines for the simulation without conduction are not shown as they lie exactly on top of the lines shown here.

SPH schemes in a general struggle to show good convergence on shock problems due to their inherent discontinuous nature. Ideal convergence for shocks with the artificial viscosity set-up used in SPHENIX is only first order (i.e. $L_1 \propto h$).

Fig. 6 shows the L_1 convergence for various fields in the Sedov–Taylor blastwave as a function of mean smoothing length. Convergence here has a best case of $L_1(v) \propto h^{1/2}$ in real terms, much

slower than the expected $L_1 \propto h^{-1}$. This is primarily due to the way that the convergence is measured; the shock front is not resolved instantaneously (i.e. there is a rise in density and velocity over some small distance, reaching the maximum value at the true position) at the same position as in the analytic solution. However, all resolution levels show an accurately placed shock front and a shock width that scales linearly with resolution (see Appendix D for more information).

5.3 Gresho–Chan vortex

The Gresho–Chan vortex (Gresho & Chan 1990) is typically used to test for the conservation of vorticity and angular momentum, and is performed here in two dimensions. Generally, it is expected that the performance of SPH on this test is more dependent on the kernel employed (see Dehnen & Aly 2012), as long as a sufficient viscosity-suppressing switch is used.

5.3.1 *Initial conditions*

The initial conditions use a 2D glass file, and treat the gas with an adiabatic index $\gamma = 5/3$, constant density $\rho_0 = 1$, in a square of side-length 1. The particles are given azimuthal velocity

$$v_\phi = \begin{cases} 5r & r < 0.2 \\ 2 - 5r & 0.2 \leq r < 0.4 \\ 0 & r \geq 0.4 \end{cases} \tag{36}$$

with the pressure set so that the system is in equilibrium as

$$P_0 = \begin{cases} 5 + 12.5r^2 & r < 0.2 \\ 9 + 12.5r^2 - 20r + 4\log(5r) & 0.2 \leq r < 0.4 \\ 3 + 4\log(2) & r \geq 0.4 \end{cases} \tag{37}$$

where $r = \sqrt{x^2 + y^2}$ is the distance from the box centre.

5.3.2 *Results*

Fig. 7 shows the state of a high resolution (using a glass containing 512^2 particles) result after one full rotation at the peak of the vortex ($r = 0.2$, $t = 1.3$). The vortex is well supported, albeit with some scatter, and the peak of the vortex is preserved. There has been some transfer of energy to the centre with a higher density and internal energy than the analytic solution due to the viscosity switch (shown on the bottom right) having a very small, but nonzero, value. This then allows for some of the kinetic energy to be transformed to thermal, which is slowly transported towards the centre as this is the region with the lowest thermal pressure.

Fig. 8 shows the convergence properties for the vortex, with the SPHENIX scheme providing convergence as good as $L_1 \propto h^{0.7}$ for the azimuthal velocity. As there are no non-linear gradients in internal energy present in the simulation, there is a very little difference between the simulations performed with and without conduction at each resolution level due to the non-activation of equation 31. This level of convergence is similar to the rate seen in Dehnen & Aly (2012) implying that the SPHENIX scheme, even with its less complex viscosity limiter, manages to recover some of the benefits of the more complex *Inviscid* scheme thanks to the novel combination of switches employed.

5.4 Noh problem

The Noh (1987) problem is known to be extremely challenging, particularly for particle-based codes, and generally requires a high

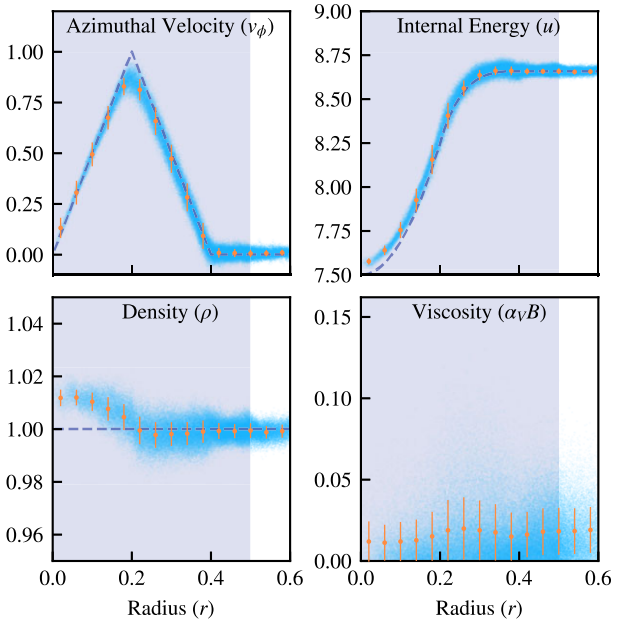

Figure 7. Gresho vortex at $t = 1.3$ after one rotation of the vortex peak with the SPHENIX scheme using a background resolution of 512^2 and with a mach number of $\mathcal{M} = 0.33$. Here, the blue points show all particles in the volume, the purple band the region used for convergence testing in Fig. 8, and the purple dashed line shows the analytic solution. The viscosity switch panel shows a very low maximal value (0.15) relative to the true maximum allowed by the code ($\alpha_V B = 2.0$), with the mean value (orange points with error bars indicating one standard deviation of scatter) of around 0.02 showing an excellent activation of the viscosity reducing switches throughout the SPHENIX scheme.

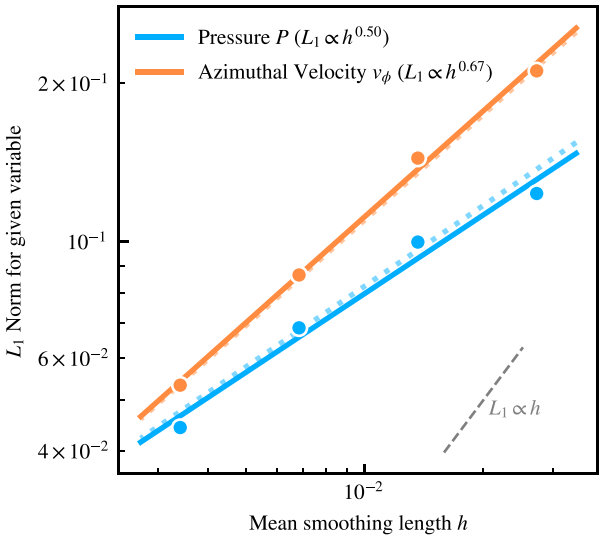

Figure 8. L_1 convergence with a mean smoothing length for various particle fields in the Gresho vortex test, measured against the analytic solution within the shaded region of Fig. 7. Each set of points shows a measured value from an individual simulation, with the lines showing a linear fit to the data in logarithmic space. The solid lines show results obtained with the full SPHENIX scheme, with dotted lines showing the results with the artificial conduction scheme disabled.

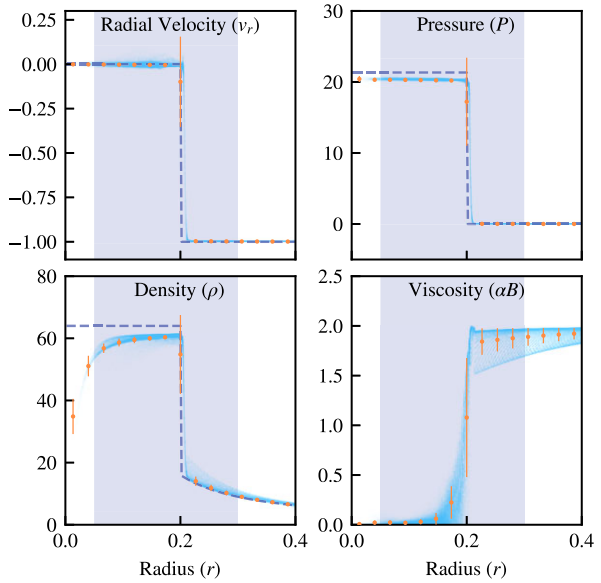

Figure 9. Noh problem simulation state at $t = 0.6$, showing a random sub-set of 1/100th of all of the particles plotted as blue points, the analytical solution as a dashed purple line, and binned quantities as orange points with error bars showing one standard deviation of scatter in that bin. The background shaded band shows the region considered for convergence in Fig. 11, with this figure showing the highest resolution simulation performed, using 512^3 particles. This simulation state is also visualized in Fig. 10.

particle number to correctly capture due to an unresolved convergence point. It tests a converging flow that results in a strong radial shock. This is an extreme, idealized, version of an accretion shock commonly present within galaxy formation simulations.

5.4.1 Initial conditions

There are many ways to generate initial conditions, from very simple schemes to schemes that attempt to highly optimize the particle distribution (see e.g. Rosswog 2020a). Here, we use a simple initial condition, employing a body-centred cubic lattice distribution of particles in a periodic box. The velocity of the particles is then set such that there is a convergent flow towards the centre of the box,

$$\vec{v} = -\frac{\vec{C} - \vec{x}}{|\vec{C} - \vec{x}|} \tag{38}$$

with $\vec{C} = 0.5L(1, 1, 1)$, where L is the box side-length, the coordinate at the centre of the volume. This gives every particle a speed of unity, meaning those in the centre will have extremely high relative velocities. We cap the minimal value of $|\vec{C} - \vec{x}|$ to be $10^{-10}L$ to prevent a singularity at small radii.

The simulation is performed in a dimensionless co-ordinate system, with a box-size of $L = 1$.

5.4.2 Results

The state of the simulation is shown at time $t = 0.6$ in Fig. 9 and visualized in Fig. 10, which shows the radial velocity, which should be zero inside of the shocked region (high density in Fig. 10), and the same as the initial conditions (i.e. -1 everywhere) elsewhere. This behaviour is captured well, with a small amount of scatter, corresponding to the small radial variations in density shown in the image.

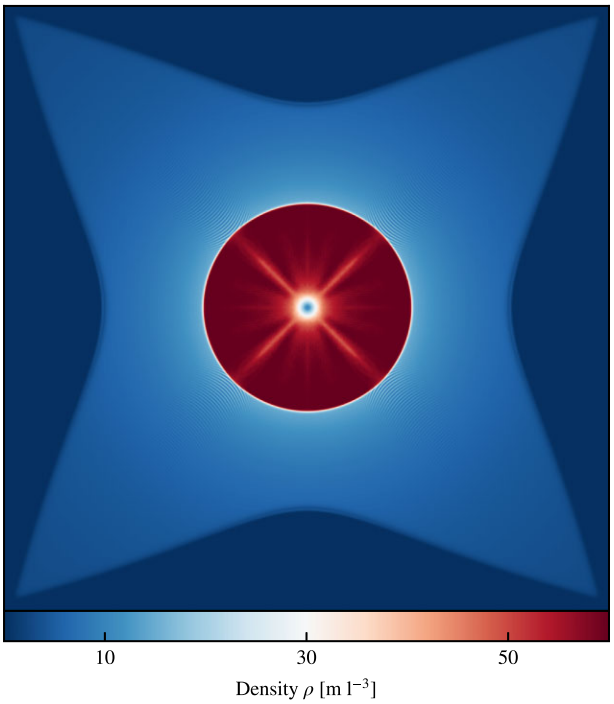

Figure 10. A density slice through the centre of the Noh problem at $t = 0.6$ corresponding to the particle distribution shown in Fig. 9. The SPHENIX scheme yields almost perfect spherical symmetry for the shock, but does not capture the expected high density in the central region, likely due to lower than required artificial conductivity (see Appendix E for more information).

The profile of density as a function of radius is however less well captured, with some small waves created by oscillations in the artificial viscosity parameter (see e.g. Rosswog 2020b, for a scheme that corrects for these errors). This can also be seen in the density slice, and is a small effect that also is possibly exacerbated by our non-perfect choice of initial conditions, but is also present in the implementation shown in Rosswog (2020a). The larger, more significant, density error is shown inside the central part of the shocked, high-density, region. This error is ever-present in SPH schemes, and is likely due to both a lack of artificial conduction in this central region (as indicated by Noh 1987, note the excess pressure in the centre caused by 'wall heating') and the unresolved point-of-flow convergence.

The Noh problem converges well using SPHENIX, with better than linear convergence for the radial velocity (Fig. 11; recall that for shocks SPH is expected to converge with $L_1 \propto h$).

This problem does not activate the artificial conduction in the SPHENIX implementation because of the presence of equation (34) reducing conductivity in highly viscous flows, as well as our somewhat conservative choice for artificial conduction coefficients (see Appendix E for more details on this topic). However, as these are necessary for the practical functioning of the SPHENIX scheme in galaxy formation simulations, and due to this test being highly artificial, this outcome presents little concern.

5.5 Square test

The square test, first presented in Saitoh & Makino (2013), is a particularly challenging test for schemes like SPHENIX that do not use a smoothed pressure in their equation of motion, as they typically lead to an artificial surface tension at contact discontinuities (the

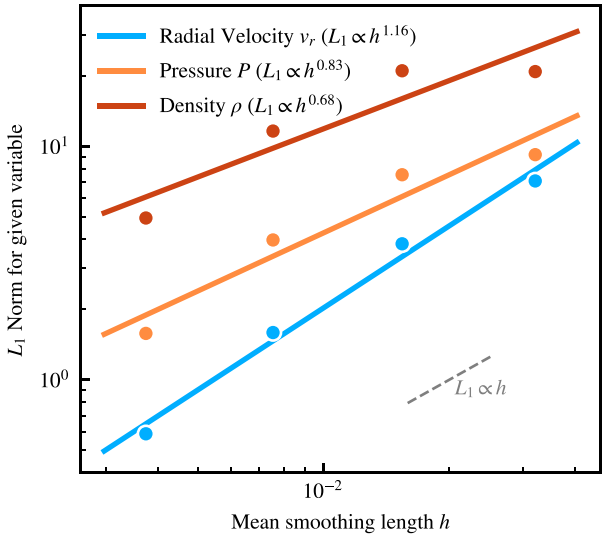

Figure 11. L_1 convergence test for various particle properties at $t = 0.6$ for the Noh problem, corresponding to the particle distribution shown in Fig. 9. The lines without conduction are not shown here as there is a little difference between the with and without conduction case, due to the extremely strong shock present in this test (leading to low values of the viscosity alpha, equation 34).

same ones that lead to the pressure blip in Section 5.1). This test is a more challenging variant of the ellipsoid test presented in Heß & Springel (2010), as the square includes sharp corners which are more challenging for conduction schemes to capture.

5.5.1 *Initial conditions*

The initial conditions are generated using equal mass particles. We set up a grid in 2D space with $n \times n$ particles, in a box of size $L = 1$. The central 0.5×0.5 square is set to have a density of $\rho_C = 4.0$, and so is replaced with a grid with $2n \times 2n$ particles, with the outer region having $\rho_O = 1.0$. The pressures are set to be equal with $P_C = P_O = 1.0$, with this enforced by setting the internal energies of the particles to their appropriate values. All particles are set to be completely stationary in the initial conditions with $\vec{v} = 0$. The initial conditions are not allowed to relax in any way.

5.5.2 *Results*

Fig. 12 shows the square test at $t = 4$ for four different resolution levels and three different variations on the SPHENIX scheme. By this time, the solutions are generally very stable. The top row shows the SPHENIX scheme without any artificial conduction enabled (this is achieved by setting $\alpha_{D,\,\mathrm{max}}$ to zero) and highlights the typical end state for a Density-Energy SPH scheme on this problem. Artificial surface tension leads to the square deforming and rounding to become more circular.

The bottom row shows the SPHENIX scheme with the artificial conduction switch removed; here, $\alpha_{D,\,\mathrm{min}}$ is set to the same value as $\alpha_{D,\,\mathrm{max}} = 1$. The artificial conduction scheme significantly reduces the rounding of the edges, with a rapid improvement as resolution increases. The rounding present here only occurs in the first few steps as the energy outside the square is transferred to the boundary region to produce a stable linear gradient in internal energy.

Finally, the central row shows the SPHENIX scheme, which gives a result indistinguishable from the maximum conduction scenario.

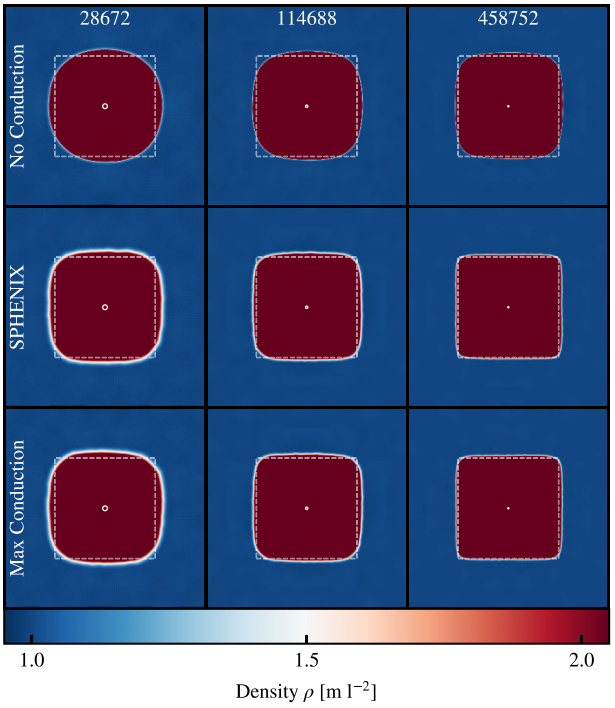

Figure 12. The density field for the square test at $t = 4$, shown at various resolution levels (different columns, numbers at the top denote the number of particles in the system) and with various modifications to the underlying SPH scheme (different rows). The dashed line shows the initial boundary of the square that would be maintained with a perfect scheme due to the uniform pressure throughout. The white circle at the centre of the square shows a typical smoothing length for this resolution level. Vertically, the scheme with no conduction is shown at the top, with the SPHENIX scheme in the middle and a scheme with the conduction coefficient set to the maximum level throughout at the bottom. The schemes with conduction maintain the square shape significantly better than the one without conduction, and the SPHENIX limiters manage to provide the appropriate amount of conduction to return to the same result as the maximum conduction case.

This is despite the initial value for the conduction coefficient $\alpha_D = 0$, meaning it must ramp up rapidly to achieve such a similar result. The SPHENIX result here shows that the choices for the conduction coefficients determined from the Sod tube (Section 5.1) are not only appropriate for that test, but apply more generally to problems that aim to capture contact discontinuities.

5.6 2D Kelvin–Helmholtz instability

The 2D Kelvin–Helmholtz instability is presented below. This test is a notable variant on the usual Kelvin–Helmholtz test as it includes a density jump at constant pressure (i.e. yet another contact discontinuity). This version of the Kelvin–Helmholtz instability is performed in two dimensions. A recent, significantly more detailed, study of Kelvin–Helmholtz instabilities within SPH is available in Tricco (2019). In this section, we focus on qualitative comparisons and how the behaviour of the instability changes with resolution within SPHENIX.

5.6.1 *Initial conditions*

The initial conditions presented here are similar to those in Price (2008), where they discuss the impacts more generally of the inclusion of artificial conduction on fluid mixing instabilities. This

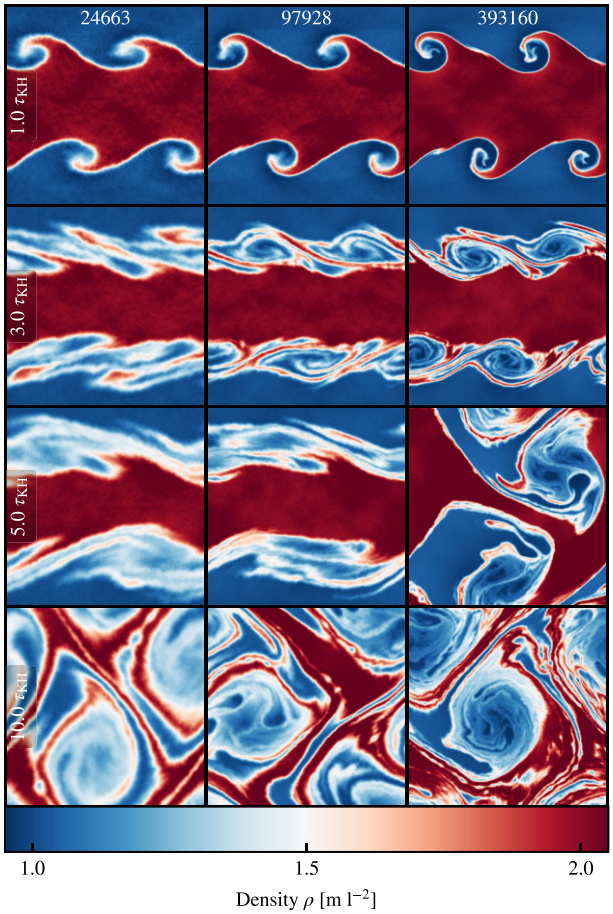

Figure 13. Density map of the standard Kelvin–Helmholtz 2D test at various resolutions (different columns, with the number of particles in the volume at the top) and at various times (different rows showing times from $t = \tau_{KH}$ to $t = 10\tau_{KH}$). Despite this being a challenging test for SPH, the instability is captured well at all resolutions, with higher resolution levels capturing finer details.

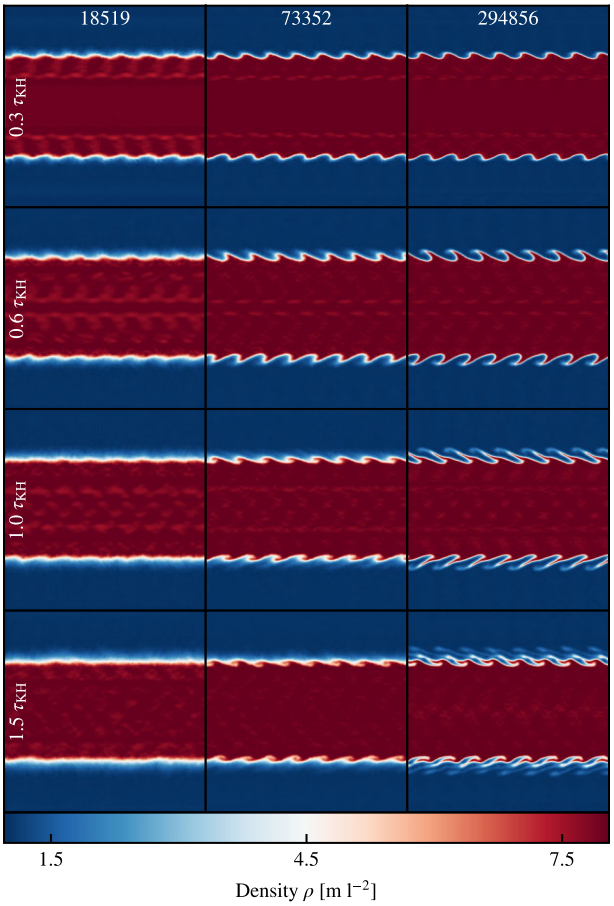

Figure 14. The same as Fig. 13, but this time using initial conditions with a significantly higher (1:8 instead of 1:2) density contrast. The initial instabilities are captured well at all resolution levels, but at the lowest level they are rapidly mixed by the artificial conduction scheme due to the lack of resolution elements in the low-density region.

is set up in a periodic box of length $L = 1$, with the central band between $0.25 < y < 0.75$ set to $\rho_C = 2$ and $v_{C,x} = 0.5$, with the outer region having $\rho_O = 1$ and $v_{O,x} = -0.5$ to set up a shear flow. The pressure $P_C = P_O = 2.5$ is enforced by setting the internal energies of the equal mass particles. Particles are initially placed on a grid with equal separations. This is the most challenging version of this test for SPH schemes to capture as it includes a perfectly sharp contact discontinuity; see Agertz et al. (2007) for more information.

We then excite a specific mode of the instability, as in typical SPH simulations un-seeded instabilities are dominated by noise and are both unpredictable and unphysical, preventing comparison between schemes.

5.6.2 Results

Fig. 13 shows the simulation after various multiples of the Kelvin–Helmholtz time-scale for the excited instability, with τ_{KH} given by

$$\tau_{KH} = \frac{(1+\chi)\lambda}{\bar{v}\sqrt{\chi}} \tag{39}$$

where $\chi = \rho_C/\rho_O = 2$ is the density contrast, $\bar{v} = v_{I,x} - v_{O,x} = 1$ the shear velocity, and $\lambda = 0.5$ the wavelength of the seed perturbation along the horizontal axis (e.g Hu et al. 2014). The figure shows three initial resolution levels, increasing from left to right. Despite this being the most challenging version of the Kelvin–Helmholtz test (at this density contrast) for a Density-Energy based SPH scheme, the instability is captured well at all resolutions, with higher resolutions allowing for more rolls of the 'swirl' to be captured. In particular, the late-time highly mixed state shows that with the conduction removed after a linear gradient in internal energy has been established, the SPHENIX scheme manages to preserve the initial contact discontinuity well. Due to the presence of explicit artificial conduction, SPHENIX seems to diffuse more than other schemes on this test (e.g. Hu et al. 2014; Wadsley et al. 2017), leading to the erausre of some small-scale secondary instabilities.

The non-linear growth rate of the swirls is resolution dependent within this test, with higher-resolution simulations showing faster growth of the largest-scale modes. This is due to better capturing of the energy initially injected to perturb the volume to produce the main instability, with higher resolutions showing lower viscous losses.

Fig. 14 shows a different initial condition where the density contrast $\chi = 8$, four times higher than the one initially presented. Because SPH is fundamentally a finite mass method, and we use equal-mass particles throughout, this is a particularly challenging test as the low-density region is resolved by so few particles. Here, we also excite an instability with a wavelength $\lambda = 0.125$, four times smaller than the one used for the $\chi = 2$ test. This value is

chosen for two reasons; it is customary to lower the wavelength of the seeded instability as the density contrast is increased when grid codes perform this test as it allows each instability to be captured with the same number of cells at a given resolution level; and also to ensure that this test is as challenging as is practical for the scheme.

SPHENIX struggles to capture the instability at very low resolutions primarily due to the lack of particles in the low-density flow (an issue also encountered by Price 2008). In the boundary region, the artificial conduction erases the small-scale instabilities on a time-scale shorter than their formation time-scale (as the boundary region is so large) and as such they cannot grow efficiently. As the resolution increases, however, SPHENIX is able to better capture the linear evolution of the instability, even capturing turn-overs and the beginning of nonlinear evolution for the highest resolution.

5.7 Blob test

The Blob test is a challenging test for SPH schemes (see Klein, McKee & Colella 1994; Springel 2005) and aims to replicate a scenario where a cold blob of gas falls through the hot IGM/CGM surrounding a galaxy. In this test, a dense sphere of cold gas is placed in a hot, low density, and supersonic wind. Ideally, the blob should breakup and dissolve into the wind, but Agertz et al. (2007) showed that the inability of traditional SPH schemes to exchange entropy between particles prevents this from occurring. The correct, specific rate at which the blob should mix with its surroundings, as well as the structure of the blob whilst it is breaking up, are unknown.

5.7.1 Initial conditions

There are many methods to set up the initial conditions for the Blob test, including some that excite an instability to ensure that the blob breaks up reliably (such as those used in Hu et al. 2014). Here, we excite no such instabilities and simply allow the simulation to proceed from a very basic particle set-up with a perfectly sharp contact discontinuity. The initial conditions are dimensionless in nature, as the problem is only specified in terms of the Mach number of the background medium and the blob density contrast.

To set up the initial particle distribution, we use two-body centred cubic lattices, one packed at a high density (for the blob, $\rho_{\rm blob} = 10$) and one at low density (for the background medium, $\rho_{\rm bg} = 1$). The low-density lattice is tiled four times in the x direction to make a box of size $4 \times 1 \times 1$, and at [0.5, 0.5, 0.5] a sphere of radius 0.1 is removed and filled in with particles from the high-density lattice. The particles in the background region are given a velocity of $v_{\rm bg} = 2.7$ (with the blob being stationary), and the internal energy of the gas everywhere is scaled such that the background medium has a mach number of $\mathcal{M} = 2.7$ and the system is in pressure equilibrium everywhere.

5.7.2 Results

The blob is shown at a number of resolution levels at various times in Fig. 15. At all resolution levels, the blob mixes well with the background medium after a few Kelvin–Helmholtz timescales (see equation 39 for how this is calculated; here we assume that the wavelength of the perturbation is the radius of the blob).[6] The rate of mixing is consistent amongst all resolution levels, implying that the artificial conduction scheme is accurately capturing unresolved mixing at lower resolutions.

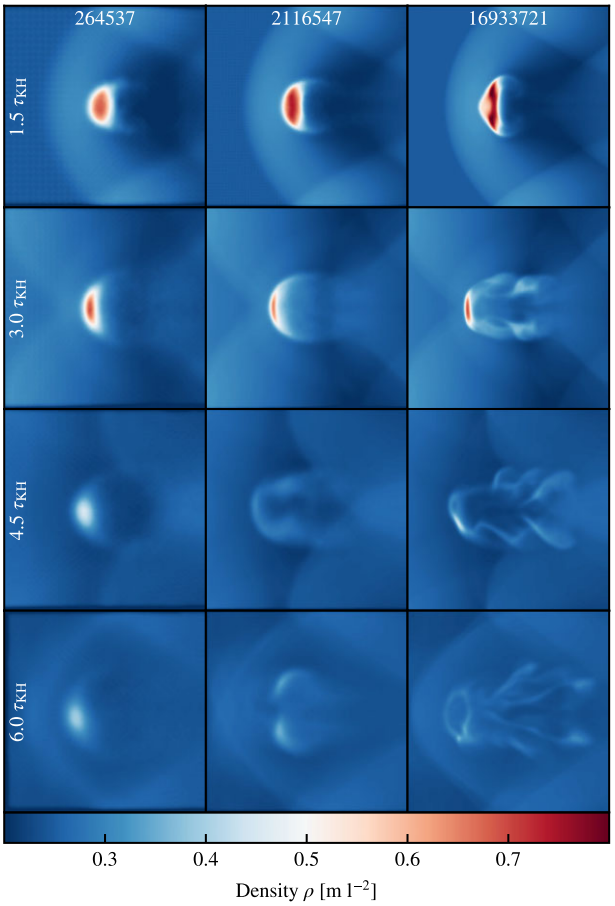

Figure 15. Time evolution of the blob within the supersonic wind at various resolution levels (different columns; the number of particles in the whole volume is noted at the top) and at various times (expressed as a function of the Kelvin–Helmholtz time for the whole blob $\tau_{\rm KH}$; different rows). The projected density is shown here to enable all layers of the 3D structure to be seen. At all resolution levels, the blob mixes with the surrounding medium (and importantly mixes phases with the surrounding medium), with higher resolution simulations displaying more thermal instabilities that promote the breaking up of the blob.

The rate of mixing of the blob is broadly consistent with that of modern SPH schemes and grid codes, however our set of initial conditions appear to mix slightly slower (taking around $\sim 4 - 6\tau_{\rm KH}$ to fully mix) than those used by other contemporary works (Agertz et al. 2007; Read & Hayfield 2012; Hu et al. 2014), possibly due to the lack of perturbation seeding (see Read et al. 2010, Appendix B for more details). When testing these initial conditions with a scheme that involves a Riemann solver or a Pressure-based scheme (see Appendix F), the rate of mixing is qualitatively similar to the one presented here. SPHENIX is unable to fully capture the crushing of the blob from the centre outwards seen in grid codes and other SPH formulations using different force expressions (Wadsley et al. 2017), rather preferring to retain a 'plate' of dense gas at the initial point of the blob that takes longer to break up. A potential explanation for this difference is some residual surface tension in SPHENIX. In these highly dynamic situations, it may not be possible for the artificial conduction to establish a smooth internal energy profile rapidly enough for small-scale instabilities on the surface to assist in the breakup of the blob.

[6]Note that here the Kelvin–Helmholtz time-scale is 1.1 times the cloud crushing time-scale (Agertz et al. 2007).

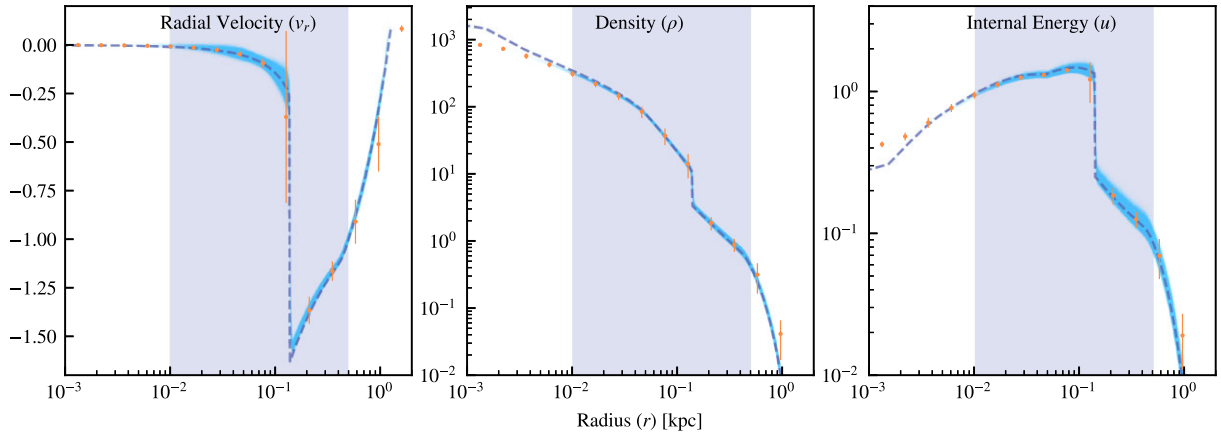

Figure 16. State of the Evrard sphere at $t = 0.8$ for a resolution of 10^7 particles. A random sub-set of 1/10th of the particles is shown in blue, with the solution from a high resolution 1D grid code shown as a purple dashed line. The orange points with error bars show the median within a radial equally log-spaced bin with the bar showing one standard deviation of scatter. The shaded band in the background shows the region considered for the convergence test in Fig. 17.

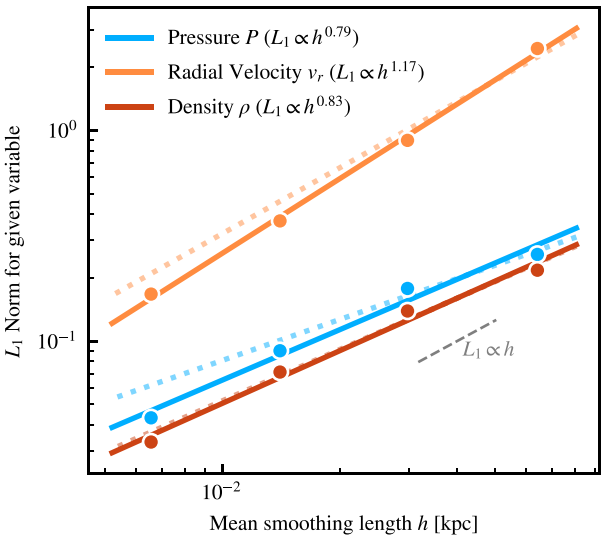

Figure 17. L_1 convergence for various gas properties for the Evrard collapse sphere at $t = 0.8$. The region considered for convergence here is the purple band shown in Fig. 16. The SPHENIX scheme is shown with the points and linear fits in solid, and the same scheme is shown with artificial conduction turned off as dotted lines. Artificial conduction significantly improves convergence here as it helps stabilize the thermal properties of the initially randomly placed particles.

At low resolutions, it is extremely challenging for the method to capture the break-up of the blob as there are very few particles in the background medium to interact with the blob due to the factor of 10 density contrast.

5.8 Evrard collapse

The Evrard collapse (Evrard 1988) test takes a large sphere of self-gravitating gas, at low energy and density, which collapses in on itself, causing an outward moving accretion shock. This test is of particular interest for cosmological and astrophysical applications as it allows for the inspection of the coupling between the gravity and hydrodynamics solver.

5.8.1 Initial conditions

Gas particles are set up in a sphere with an adiabatic index of $\gamma = 5/3$, sphere mass $M = 1$, sphere radius $R = 1$, initial density profile $\rho(r) = 1/2\pi r$, and in a very cold state with $u = 0.05$, with the gravitational constant $G = 1$. These initial conditions are created in a box of size 100, ensuring that effects from the periodic boundary are negligible. Unfortunately, due to the non-uniform density profile, it is considerably more challenging to provide relaxed initial conditions (or use a glass file). Here, positions are simply drawn randomly to produce the required density profile.

The Evrard collapse was performed at four resolution levels, with total particle numbers in the sphere being 10^4, 10^5, 10^6, and 10^7. The gravitational softening was fixed at 0.001 for the 10^6 resolution level, and this was scaled with $m^{-1/3}$ with m being the particle mass for the other resolution levels. The simulations were performed once with artificial conduction enabled (the full SPHENIX scheme), and once with it disabled.

5.8.2 Results

The highest resolution result (10^7 particles) with the full SPHENIX scheme is shown in Fig. 16. This is compared against a high resolution grid code[7] simulation performed in 1D, and here SPHENIX shows an excellent match to the reference solution. The shock at around $r = 10^{-1}$ is sharply resolved in all variables, and the density and velocity profiles show excellent agreement. In the centre of the sphere, there is a slight deviation from the reference solution for the internal energy and density (balanced to accurately capture the pressure in this region) that remains even in the simulation performed without artificial conduction (omitted for brevity, as the simulation without conduction shows similar results to the simulation with conduction, with the exception of the conduction reducing scatter in the internal energy profile). This is believed to be an artefact of the initial conditions, however it was not remedied by performing simulations at higher resolutions.

The convergence properties of the Evrard sphere are shown in Fig. 17. The velocity profile shows a particularly excellent result

[7] `HydroCode1D`, see https://github.com/bwvdnbro/HydroCode1D and the SWIFT repository for more details.

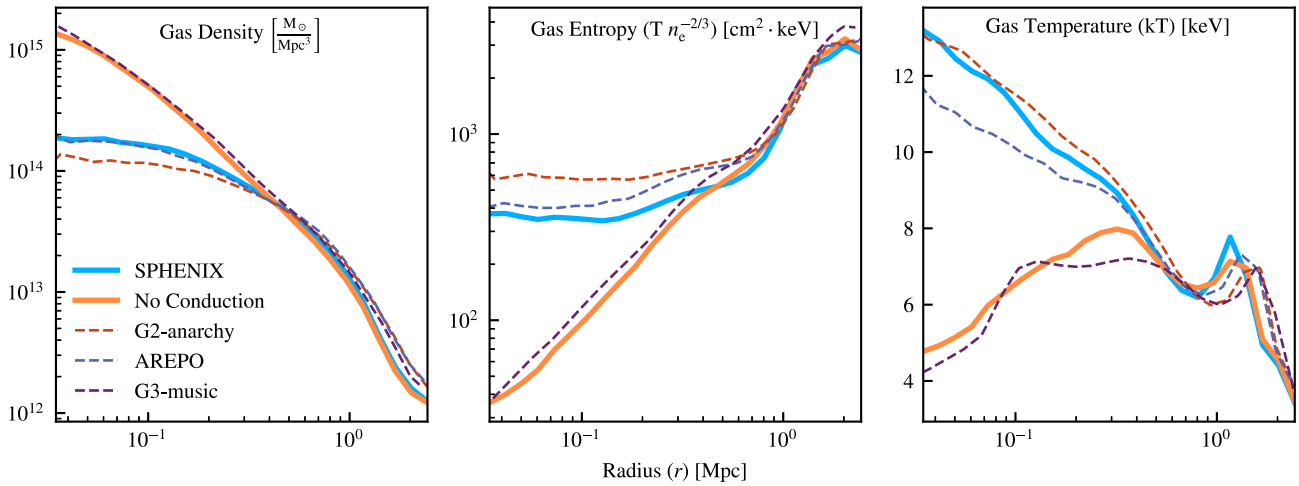

Figure 18. Thermodynamics profiles for the nIFTy cluster at $z = 0$ with five codes and schemes. The solid lines show those simulated with SWIFT, with the blue line showing the full SPHENIX scheme, and the orange line showing SPHENIX without artificial conduction. The dashed lines were extracted directly from Sembolini et al. (2016) and show a modern Pressure-Entropy scheme (G2-anarchy; Schaye et al. 2015, appendix A), a moving mesh finite volume scheme (AREPO; Springel 2010), and a traditional SPH scheme (G3-music; Springel 2005).

with greater than linear convergence demonstrated. The thermodynamic properties show roughly linear convergence. Of particular note is the difference between the convergence properties of the simulations with and without artificial conduction; those with this feature of SPHENIX enabled converge at a more rapid rate. This is primarily due to the stabilizing effect of the conduction on the internal energy profile. As the particles are initially placed randomly, there is some scatter in the local density field at all radii. This is quickly removed by adiabatic expansion in favour of scatter in the internal energy profile, which can be stabilized by the artificial conduction.

5.9 nIFTy cluster

The nIFTy cluster comparison project, Sembolini et al. (2016), uses a (non-radiative, cosmological) cluster-zoom simulation to evaluate the efficacy of various hydrodynamics and gravity solvers. The original paper compared various types of schemes, from traditional SPH (Gadget; Springel 2005) to a finite volume adaptive mesh refinement scheme (RAMSES; Teyssier 2002). The true answer for this simulation is unknown, but it is a useful case to study the different characteristics of various hydrodynamics solvers.

In Fig. 18, the SPHENIX scheme is shown with and without artificial conduction against three reference schemes from Sembolini et al. (2016). Here, the centre of the cluster was found using the VELOCIraptor (Elahi et al. 2019) friends-of-friends halo finder, and the particle with the minimum gravitational potential was used as the reference point.

The gas density profile was created using 25 equally log-spaced radial bins, with the density calculated as the sum of the mass within a shell divided by the shell volume. SPHENIX scheme shows a similar low-density core as AREPO, with the no conduction scheme resulting in a cored density profile similar to the traditional SPH scheme from Sembolini et al. (2016).

The central panel of Fig. 18 shows the 'entropy' profile of the cluster; this is calculated as $Tn_e^{-2/3}$ with n_e the electron density (assuming primordial gas, this is $n_e = 0.875\rho/m_H$ with m_H the mass of a hydrogen atom) and T the gas temperature. Each was calculated individually in the same equally log-spaced bins as the density profile, with the temperature calculated as the mass-weighted temperature within that shell. The rightmost panel shows this mass-weighted temperature profile, with SPHENIX showing slightly higher temperatures in the central region than AREPO, matching G2-anarchy instead. This high-temperature central region, along with a low-density centre, lead to the 'cored' (i.e. flat, with high values of entropy, at small radii) entropy profile for SPHENIX.

The cored central entropy profile with SPHENIX is attained primarily due to the artificial conduction scheme and is not due to the other improvements over the traditional SPH base scheme (including for example the artificial viscosity implementation). We note again that there was no attempt to calibrate the artificial conduction scheme to attain this result on the nIFTy cluster, and any and all parameter choices were made solely based on the Sod shock tube in Section 5.1.

In Fig. 19, a projected mass-weighted temperature image of the cluster is shown. The image demonstrates how the artificial conduction present in the SPHENIX scheme promotes phase mixing, resulting in the cored entropy profile demonstrated in Fig. 18.

The temperature distribution in the SPH simulation without conduction appears noisier, due to particles with drastically different phases being present within the same kernel. This shows how artificial conduction can lead to sharper shock capture as the particle distribution is less susceptible to this noise, enabling a cleaner energy transition between the pre- and post-shock region.

6 CONCLUSIONS

We have presented the SPHENIX SPH scheme and its performance on seven hydrodynamics tests. The scheme has been demonstrated to show convergent (with resolution) behaviour on all these tests. In summary:

(i) SPHENIX is an SPH scheme that uses Density-Energy SPH as a base, with added artificial viscosity for shock capturing and artificial conduction to reduce errors at contact discontinuities and to promote phase mixing.

(ii) A novel artificial conduction limiter allows SPHENIX to be used with energy injection feedback schemes (such as those used in

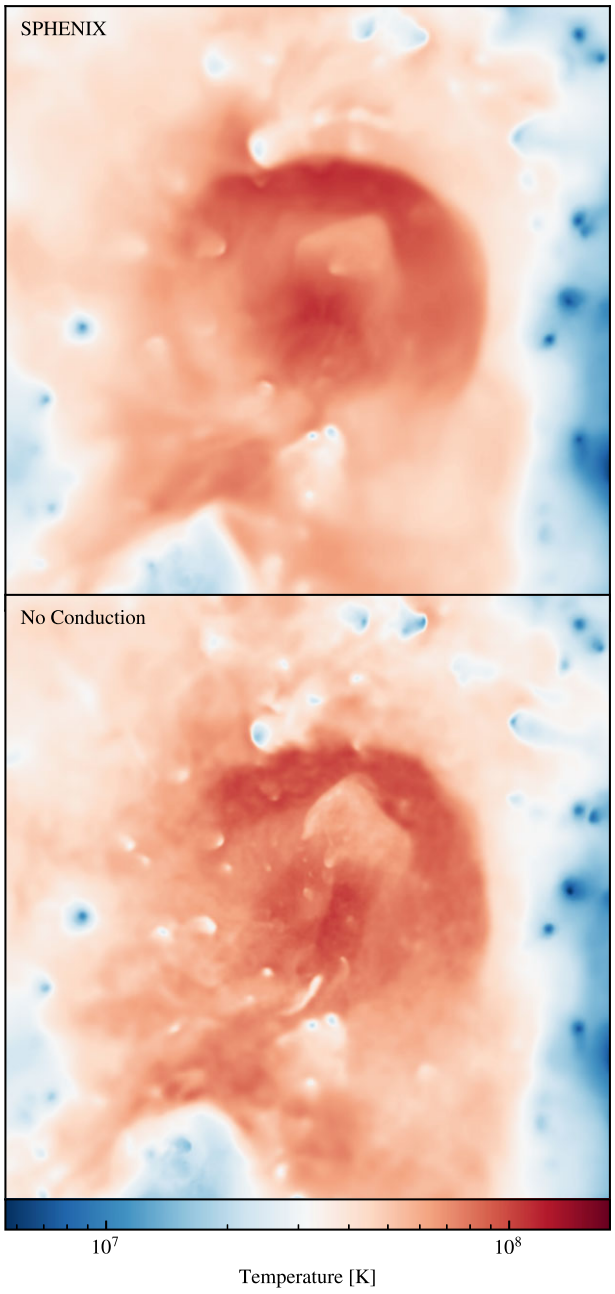

Figure 19. Image of the nIFTY cluster, as a projected mass-weighted temperature map, shown for the SPHENIX scheme with (top) and without artificial conduction enabled (bottom). The image shows a 5 Mpc wide view, centred on the most bound particle in the halo.

EAGLE) by reducing conduction across shocks and other regions where the artificial viscosity is activated.

(iii) The artificial viscosity and conduction scheme coefficients were determined by ensuring good performance on the Sod Shock tube test, and remain fixed for all other tests.

(iv) The modified Inviscid SPH (Cullen & Dehnen 2010) scheme captures strong shocks well, ensuring energy conservation, as shown by the Sedov-Taylor blastwave test, but the smooth nature of SPH prevents rapid convergence with resolution.

(v) The use of the Balsara (1989) switch in SPHENIX was shown to be adequate to ensure that the Gresho–Chan vortex is stable. Convergence on this test was shown to be faster than in Cullen & Dehnen (2010).

(vi) The artificial conduction scheme was shown to work adequately to capture thermal instabilities in both the Kelvin-Helmholtz and Blob tests, with contact discontinuities well preserved when required.

(vii) SPHENIX performed well on both the Evrard collapse and nIFTY cluster problems, showing that it can couple to the FMM gravity solver in SWIFT and that the artificial conduction scheme can allow for entropy cores in clusters.

(viii) SPHENIX is implemented in the SWIFT code and is available fully open source to the community.

SPHENIX hence achieves its design goals; the Lagrangian nature of the scheme allows for excellent coupling with gravity; the artificial conduction limiter allows the injection of energy as in the EAGLE sub-grid physics model; and the low cost-per-particle and lack of matrices carried on a particle-by-particle basis provide for a very limited computational cost (see Borrow et al. 2019, for a comparison of computational costs between a scheme like SPHENIX and the GIZMO-like schemes also present in SWIFT).

SOFTWARE CITATIONS

This paper made use of the following software packages:

(i) SWIFT (Schaller et al. 2018)
(ii) PYTHON (van Rossum & Drake Jr 1995), with the following libraries

(a) NUMPY (Harris et al. 2020)
(b) SCIPY (SciPy 1.0 Contributors et al. 2020)
(c) NUMBA (Lam, Pitrou & Seibert 2015)
(d) MATPLOTLIB (Hunter 2007)
(e) SWIFTSIMIO (Borrow & Borrisov 2020)

ACKNOWLEDGEMENTS

The authors thank Folkert Nobels for providing initial conditions for Appendix C, and the anonymous referee for their comments that improved the paper. JB is supported by STFC studentship ST/R504725/1. MS is supported by the Netherlands Organisation for Scientific Research (NWO) through VENI grant 639.041.749. RGB is supported by the Science and Technology Facilities Council ST/P000541/1. This work used the DiRAC@Durham facility managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). The equipment was funded by BEIS capital funding via STFC capital grants ST/K00042X/1, ST/P002293/1, ST/R002371/1, and ST/S002502/1, Durham University, and STFC operations grant ST/R000832/1. DiRAC is part of the National e-Infrastructure.

DATA AVAILABILITY

All code and initial conditions used to generate the simulations is open source as part of SWIFT version 0.9.0 (Schaller et al. 2018).[8] As the simulations presented in this paper are small, test, simulations that can easily be repeated, the data is not made immediately available.

[8]Available from http://swift.dur.ac.uk, https://gitlab.cosma.dur.ac.uk/swift/swiftsim, and https://github.com/swiftsim/swiftsim.

REFERENCES

Agertz O. et al., 2007, MNRAS, 380, 963
Balsara D. S., 1989, PhD thesis, Univ. Illinois at Urbana-Champaign
Bauer A., Springel V., 2012, MNRAS, 423, 2558
Benz W., 1988, Comput. Phys. Commun., 48, 97
Borrow J., Borrisov A., 2020, J. Open Sour. Softw., 5, 2430
Borrow J., Bower R. G., Draper P. W., Gonnet P., Schaller M., 2018, Proceedings of the 13th SPHERIC International Workshop, June 26–28 2018. National University of Ireland, Galway, Ireland, p. 44
Borrow J., Vandenbroucke B., Schaller M., 2019, in Proceedings of the 14th SPHERIC International Workshop, June 25–27 2019. University of Exeter, United Kingdom
Borrow J., Schaller M., Bower R. G., 2021, MNRAS, 505, 2316
Cen R., 1992, ApJS, 78, 341
Cen R., Ostriker J. P., 1992, ApJ, 399, L113
Courant R., Friedrichs K., Lewy H., 1928, Math. Ann., 100, 32
Crain R. A., Eke V. R., Frenk C. S., Jenkins A., McCarthy I. G., Navarro J. F., Pearce F. R., 2007, MNRAS, 377, 41
Cullen L., Dehnen W., 2010, MNRAS, 408, 669
Dalla Vecchia C., Schaye J., 2008, MNRAS, 387, 1431
Dalla Vecchia C., Schaye J., 2012, MNRAS, 426, 140
Debackere S. N. B., Schaye J., Hoekstra H., 2020, MNRAS, 492, 2285
Dehnen W., 2014, Computat. Astrophys. Cosmol., 1, 1
Dehnen W., Aly H., 2012, MNRAS, 425, 1068
Dolag K., Jubelgas M., Springel V., Borgani S., Rasia E., 2004, ApJ, 606, L97
Durier F., Dalla Vecchia C., 2012, MNRAS, 419, 465
Elahi P. J., Cañas R., Poulton R. J. J., Tobar R. J., Willis J. S., Lagos C. d. P., Power C., Robotham A. S. G., 2019, Publ. Astron. Soc. Aust., 36, e021
Ettori S., Dolag K., Borgani S., Murante G., 2006, MNRAS, 365, 1021
Evrard A. E., 1988, MNRAS, 235, 911
Gingold R. A., Monaghan J. J., 1977, MNRAS, 181, 375
Greengard L., Rokhlin V., 1987, J. Comput. Phys., 73, 325
Gresho P. M., Chan S. T., 1990, Int. J. Numer. Methods Fluids, 11, 621
Gresho P. M., Sani R. L., 1990, Int. J. Numer. Methods Fluids, 11, 951
Guillet T., Pakmor R., Springel V., Chandrashekar P., Klingenberg C., 2019, MNRAS, 485, 4209
Harris C. R. et al., 2020, Nature, 585, 357
Hernquist L., Katz N., 1989, ApJS, 70, 419
Heß S., Springel V., 2010, MNRAS, 406, 2289
Hopkins P. F., 2013, MNRAS, 428, 2840
Hopkins P. F., 2015, MNRAS, 450, 53
Hu C.-Y., Naab T., Walch S., Moster B. P., Oser L., 2014, MNRAS, 443, 1173
Hunter J. D., 2007, Comput. Sci. Eng., 9, 90
Klein R. I., McKee C. F., Colella P., 1994, ApJ, 420, 213
Lam S. K., Pitrou A., Seibert S., 2015, in Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC. LLVM '15. Association for Computing Machinery, New York, NY, USA doi:10.1145/2833157.2833162
Lattanzio J. C., Monaghan J. J., Pongracic H., Schwarz M. P., 1986, SIAM J. Sci. Stat. Comput., 7, 591
Lucy L. B., 1977, AJ, 82, 1013
Monaghan J. J., 1992, ARA&A, 30, 543
Monaghan J., 1997, J. Comput. Phys., 136, 298
Monaghan J. J., 2002, MNRAS, 335, 843
Monaghan J., Gingold R., 1983, J. Comput. Phys., 52, 374
Morris J., Monaghan J., 1997, J. Comput. Phys., 136, 41
Navarro J. F., White S. D. M., 1993, MNRAS, 265, 271
Navarro J. F., Frenk C. S., White S. D. M., 1996, ApJ, 462, 563
Noh W. F., 1987, J. Comput. Phys., 72, 78
Oppenheimer B. D., Schaye J., Crain R. A., Werk J. K., Richings A. J., 2018, MNRAS, 481, 835
Ploeckinger S., Schaye J., 2020, MNRAS, 497, 4857
Price D. J., 2007, Publ. Astron. Soc. Aust., 24, 159
Price D. J., 2008, J. Comput. Phys., 227, 10040
Price D. J., 2012, J. Comput. Phys., 231, 759
Price D. J. et al., 2018, Publ. Astron. Soc. Aust., 35, e031
Read J. I., Hayfield T., 2012, MNRAS, 422, 3037
Read J. I., Hayfield T., Agertz O., 2010, MNRAS, 405, 1513
Revaz Y., Jablonka P., 2012, A&A, 538, A82
Rosswog S., 2020a, MNRAS, 498, 4230
Rosswog S., 2020b, ApJ, 898, 60
Saitoh T. R., Makino J., 2009, ApJ, 697, L99
Saitoh T. R., Makino J., 2013, ApJ, 768, 44
Schaller M., Dalla Vecchia C., Schaye J., Bower R. G., Theuns T., Crain R. A., Furlong M., McCarthy I. G., 2015, MNRAS, 454, 2277
Schaller M., Gonnet P., Chalk A. B. G., Draper P. W., 2016, Proceedings of the Platform for Advanced Scientific Computing Conference on - PASC '16. Lausanne, Switzerland, p. 2:1–2:10
Schaller M., Gonnet P., Draper P. W., Chalk A. B. G., Bower R. G., Willis J., Hausammann L., 2018, Astrophysics Source Code Library, record ascl:1805.020
Schaye J., Dalla Vecchia C., 2008, MNRAS, 383, 1210
Schaye J. et al., 2015, MNRAS, 446, 521
Contributors SciPy 1.0, 2020, Nature Methods, 17, 261
Sedov, 1959, Similarity and Dimensional Methods in Mechanics - NASA/ADS.
Sembolini F. et al., 2016, MNRAS, 457, 4063
Sod G. A., 1978, J. Comput. Phys., 27, 1
Springel V., 2005, MNRAS, 364, 1105
Springel V., 2010, MNRAS, 401, 791
Springel V., Hernquist L., 2002, MNRAS, 333, 649
Springel V., Hernquist L., 2003, MNRAS, 339, 289
Stern J., Fielding D., Faucher-Giguère C.-A., Quataert E., 2019, MNRAS, 488, 2549
Taylor G., 1950, Proc. R. Soc. London Ser. A, 201, 159
Teyssier R., 2002, A&A, 385, 337
Tricco T. S., 2019, MNRAS, 488, 5210
van Rossum G., Drake F. L. Jr, 1995, in Vol. 620, Python Tutorial. Centrum voor Wiskunde en Informatica, Amsterdam
Vila J. P., 1999, Math. Models Methods Appl. Sci., 9, 161
Wadsley J., Stadel J., Quinn T., 2004, New Astron., 9, 137
Wadsley J. W., Veeravalli G., Couchman H. M. P., 2008, MNRAS, 387, 427
Wadsley J. W., Keller B. W., Quinn T. R., 2017, MNRAS, 471, 2357
Weinberger R., Springel V., Pakmor R., 2020, ApJS, 248, 32
Wendland H., 1995, Adv. Comput. Math., 4, 389
Wiersma R. P. C., Schaye J., Theuns T., Dalla Vecchia C., Tornatore L., 2009, MNRAS, 399, 574
Wood D., 1981, MNRAS, 194, 201

APPENDIX A: PARTICLE COSTS

Different SPH models require different information stored per particle. Compared to a basic, 'Traditional' SPH model, SPHENIX requires a small amount of extra data to store things like the particle-carried artificial conduction and viscosity coefficients. The amount of data required increases for more complex models, such as those making use of the full shear tensor, like Wadsley et al. (2017), or additional corrections, like Rosswog (2020b). SPH models using an Arbitrary Lagrangian-Eulerian (ALE) framework (see Vila 1999) require even more information as the particles carry flux information for use in the Riemann solver.

The amount of data required to store a single element in memory is of upmost importance when considering the speed at which a simulation will run. SPH codes, and SWIFT in particular, are bound by the memory bandwidth available, rather than the costs associated with direct computation. This means any increase in particle cost corresponds to a linear increase in the required computing time for simulation; this is why keeping the particles lean is a key requirement of the SPHENIX model. Additionally, in large simulations performed over many nodes, the bandwidth of the interconnect can further

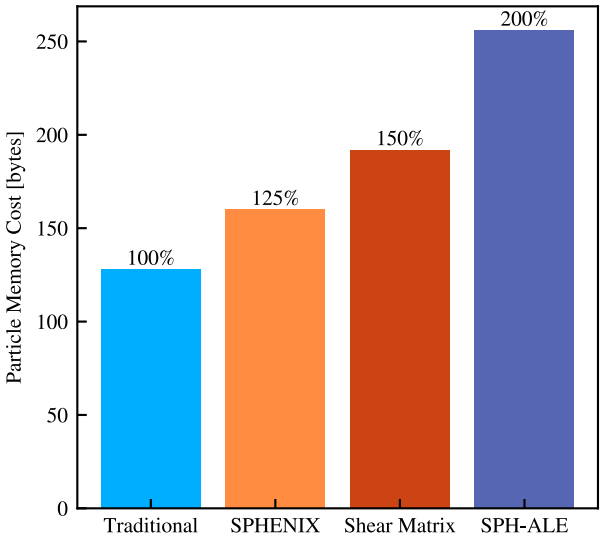

Figure A1. Cost per particle (in bytes) for four different hydrodynamics models (see text for details). Percentages are give relative to the Traditional (similar to GADGET-2, with no artificial conduction and fixed artificial viscosity coefficients) model.

become a limitation and hence keeping the memory cost of particles low is again beneficial.

In Fig. A1, we show the memory cost of four models: Traditional SPH (similar to to the one implemented in GADGET-2; Springel 2005), SPHENIX, a model with the full shear matrix, and a SPH-ALE model similar to GIZMO-MFM (Hopkins 2015), all implemented in the SWIFT framework. We see that SPHENIX only represents a 25 per cent increase in memory cost per particle for significant improvement over the traditional model.

APPENDIX B: CONDUCTION SPEED

The conduction speed (equation 26) in SPHENIX was primarily selected for numerical reasons. In a density-based scheme, it is common to see significant errors around contact discontinuities where there are large changes in density and internal energy simultaneously to produce a uniform pressure field. In Fig. 3, we demonstrated the performance of the SPHENIX scheme on one of these discontinuities, present in the Sod Shock.

In Fig. B1, we zoom in on the contact discontinuity, this time using glass files for the base initial conditions (of resolution 32^3 and 64^3), allowing for a more even distribution of particles along the horizontal axis. We use five different models,

(i) SPHENIX, the full SPHENIX model using the conduction speed from equation (26).
(ii) Pressure term, which uses only the pressure-based term from equation (26).
(iii) Velocity term, which only uses the velocity-based term from equation (26).
(iv) No conduction, which sets the conduction speed to zero.
(v) Max conduction, which sets α_D to unity everywhere, and uses the conduction speed from equation (26).

The first thing to note here is that the pressure term provides the vast majority of the conduction speed, highlighting the importance of this form of bulk conduction in SPHENIX and other models employing a density-based equation of motion. Importantly, the conduction allows for the 'pressure blip' to be reduced to a level where there is no longer a discontinuity in pressure (i.e. there is a smooth gradient with x). Although the velocity term is able to marginally reduce the size of the blip relative to the case without conduction, it is unable to fully stabilize the solution alone. Pressure blips can lead to large pressure differences between individual particles, then leading to the generation of a divergent flow around the point where the contact discontinuity resides. This is the primary motivation for the inclusion of the velocity divergence-based term in the conduction speed. Along with the conduction limiter (see equation 28 for the source term), if there is a large discontinuity in internal energy that is generating a divergent flow (and not one that is expected to do so, such as a shock), the velocity-dependent term can correct for this and smooth out the internal energy until the source of divergence disappears.

B1 Alternative conduction speeds

The SPHENIX conduction speed (equation 26) contains two components: one based on pressure differences and one based on a velocity component. In SPHENIX, as in a number of other models, this velocity component really encodes compression or expansion along the axis between particles.

The motivation for some alternative schemes (e.g. those presented in Wadsley et al. 2008, 2017) is shear between particles. To test if we see significant differences in our tests, we formulate a new conduction speed,

$$v_{D,ij} = \frac{\alpha_{D,ij}}{2}\left(\frac{|\vec{v}_{ij}\times\vec{x}_{ij}|}{|\vec{x}_{ij}|} + \sqrt{2\frac{|P_i - P_j|}{\hat{\rho}_j + \hat{\rho}_j}}\right), \tag{B1}$$

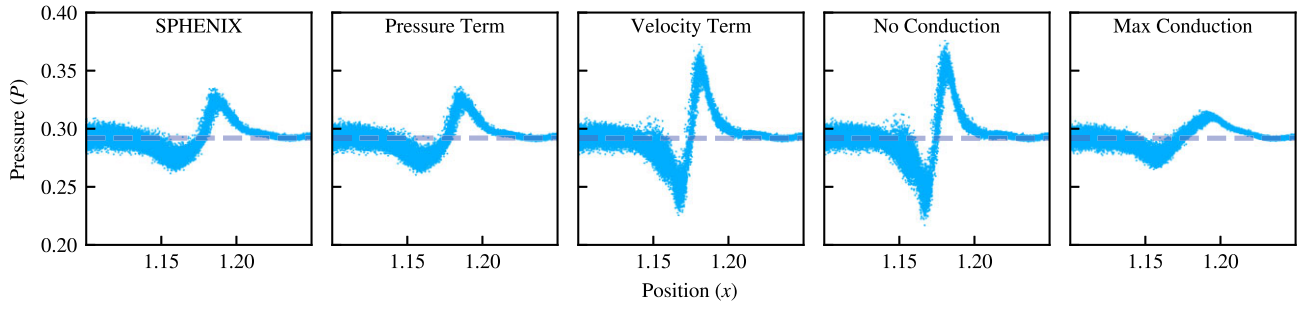

Figure B1. The pressure contact discontinuity in the Sod Shock (Fig. 3) at a resolution of 32^3 and 64^3, but using glass files instead of the Body Centered Cubic (BCC) lattices (this leads to significantly increased particle disorder, but more evenly distributes particles in the x direction enabling this figure to be clearer). Here, we show a zoomed-in representation of all particles (blue points) against the analytical solution (purple dashed line). Each sub-figure shows the simulation at the same time $t = 0.2$, but with different forms for the conduction velocity (see text for details).

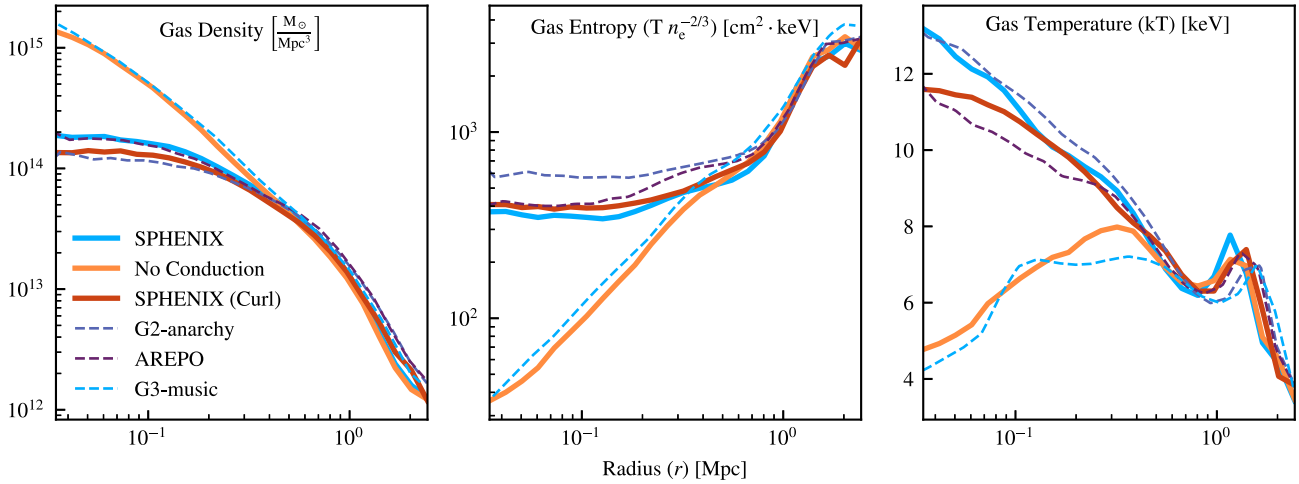

Figure B2. Reproduction of Fig. 18 but including the line (red) for the version of SPHENIX performed with an explicit shear component in the conduction speed. We see no qualitative differences between the two models, with them both providing an entropy core at a similar level.

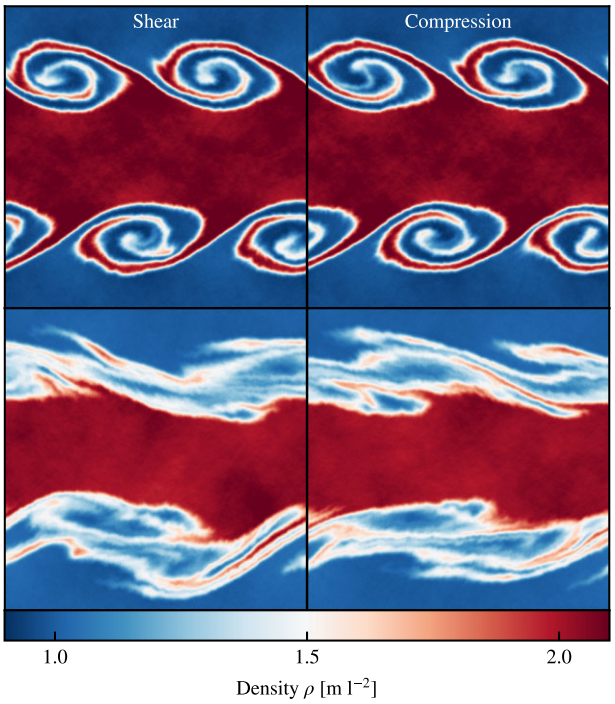

Figure B3. Kelvin–Helmholtz test with a density contrast of $\rho_C = 2$ as in Fig. 13, shown at $t = 2\tau_{KH}$ (top) and $t = 4\tau_{KH}$ (bottom). We show on the left the simulation with the shear-based conduction speed, and again the compression-based speed on the right. No significant qualitative differences are seen between the two models.

which focuses on capturing the shear component of the velocity between two particles.

We again test this new formulation on some of our example problems. First, the nIFTy cluster, presented in Fig. B2, shows little difference between the two formulations, with both providing a solution similar to other modern SPH schemes and grid codes.

The Kelvin–Helmholtz test again shows little difference (Fig. B3), although there is a slightly increased growth rate of the perturbation at late times for the shear formulation.

We find no discernible difference between the two formulations on the blob test, as this is mainly limited by the choice of Density-SPH as the base equation of motion to correctly capture the initial breakup of the blob from the centre outwards.

APPENDIX C: MAINTENANCE OF HYDROSTATIC BALANCE

The form of the conduction speed used in SPHENIX based on pressure differences (equation 26) has been conjectured to not allow for the maintenance of a pressure gradient against some external body force (for example a halo in hydrostatic equilibrium; Sembolini et al. 2016). The main concern here is that the pressure difference form of the conduction speed may allow thermal energy to travel down into the gravitational potential, heating the central regions of the halo. As SPHENIX uses an additional limiter (equation 28 for the source term) that only activates conduction in regions where the internal energy gradient cannot be represented by SPH anyway, this may be less of a concern. Additionally, there will be no conduction across accretion shocks due to the limiter in equation (33).

In Fig. C1, we show an idealized simulation of an adiabatic halo with an NFW (Navarro, Frenk & White 1996) dark matter density profile, and gas in hydrostatic equilibrium. The halo uses a fixed NFW potential in the background, with a mass of 10^{13} $M_\odot$, concentration 7.2, and a stellar bulge fraction of 1 per cent. The halo has a gas mass of 1.7×10^{12} $M_\odot$, resolved by 1067 689 particles with varying mass from 10^5 to 1.7×10^{12} $M_\odot$ with the highest resolution in the centre.

The gas in the halo is set up to be isothermal, following (Stern et al. 2019),

$$\frac{\mathrm{d}\ln P}{\mathrm{d}\ln r} = -\gamma \frac{v_c^2}{c_s^2} \tag{C1}$$

where v_c is the circular velocity of the halo. The condition used to set the initial temperatures is $v_c = c_s$, and to get the correct normalization for pressure and density the gas fraction at $R_{500,\,\mathrm{crit}}$ is used following Debackere, Schaye & Hoekstra (2020).

Fig. C1 shows that there is little difference between the result with conduction, and without it. There is a small offset in the centre where the simulation with conduction has a slightly higher energy and slightly lower density, giving a very small overall offset in pressure. This figure is shown at $t = 5$ Gyr, much longer than

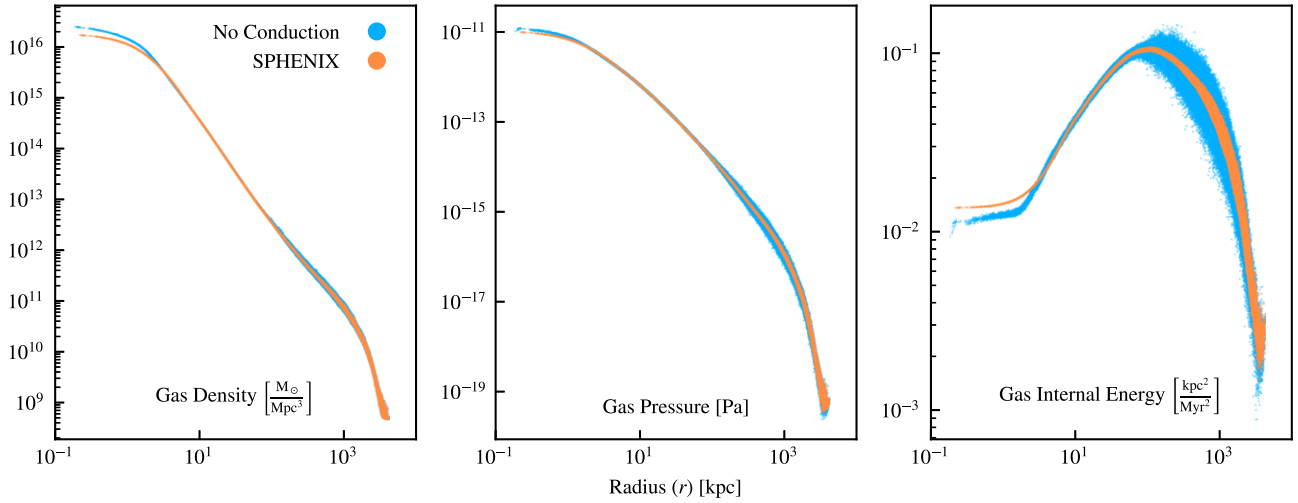

Figure C1. Profiles of the idealized NFW halo (of mass $\approx 10^{13}$ $M_\odot$, at a gas particle mass resolution of 10^5 $M_\odot$) at $t = 5$ Gyr after the initial state. Blue points show every particle in the simulation without artificial conduction enabled, with orange showing the simulation with conduction enabled. Here, the conduction can allow for a reduction in the scatter in internal energy without leading to significant conduction into the centre. The offset seen in the centre of about a factor of 1.5× originates from the smoothing of the kink around ≈0.7 kpc during the initial settling of the halo, and remained stable from that point at around $\approx t =$ 0.5 Gyr until the end of the simulation.

any realistic cluster of a similar mass would go without accretion or some other external force perturbing the pressure profile anyway. Finally, the conduction allows the noisy internal energy distribution (and additionally density distribution) to be normalized over time thanks to the inclusion of the pressure differencing term.

APPENDIX D: SEDOV BLAST

In Fig. 6, we presented the convergence properties of the Sedov blast with the SPHENIX scheme. The scheme only demonstrated convergence as $L_1(v) \propto h^{\sim 0.5}$, which is much slower than the expected convergence rate of $L_1 \propto h^1$ for shock fronts in SPH (that is demonstrated and exceeded in the Noh problem in Fig. 11). This is, however, simply an artefact of the way that the convergence is measured.

In Fig. D1, we show the actual density profiles of the shock front, by resolution (increasing as the subfigures go to the right). Note here, that the width of the shock front (from the particle distribution to the right of the vertical line to the vertical line in the analytical solution) does converge at the expected rate of $L_1 \propto 1/n^{1/3} \propto h$ with n the number of particles in the volume (in 3D).

The Sedov blast, unlike the Noh problem and Sod tubes, does not aim to reproduce a simple step function in density and velocity, but also a complex, expanding, post-shock region. The L_1 convergence is measured 'vertically' in this figure, but it is clear here that the vertical deviation from the analytical solution is not representative of the 'error' in the properties of a given particle, or in the width of the shock front. Small deviations in the position of the given particle could result in changes of orders of magnitude in the value of the L_1 norm measured for it.

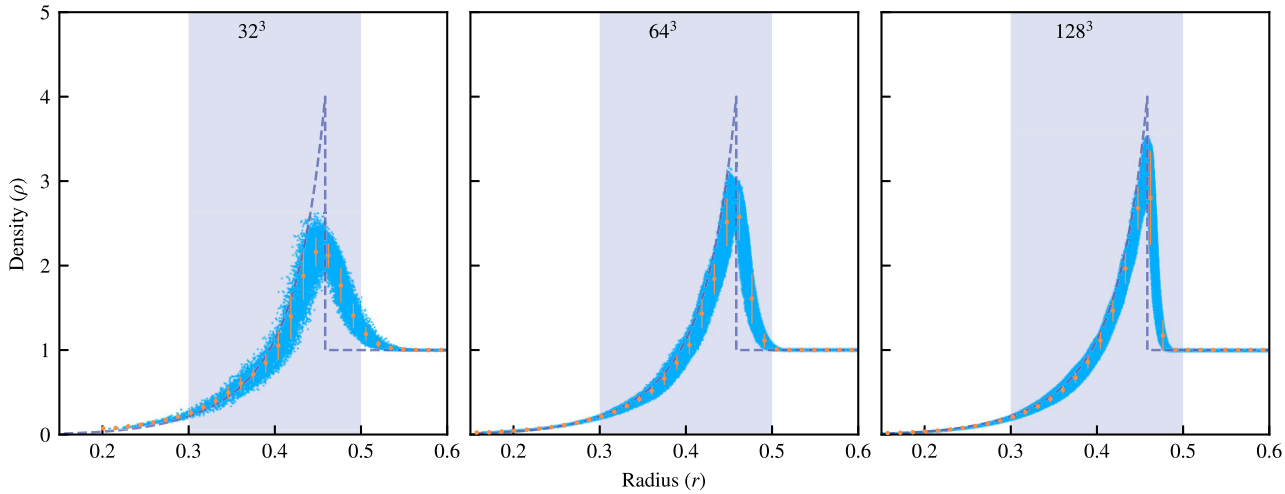

Figure D1. The density profile of the Sedov blasts initially presented in Fig. 6. The blue points show the positions of every particle in the volume, the purple dashed line the analytical prediction, and the orange points binned means with error bars showing one standard deviation. The shaded band is the region over which the convergence properties were measured. The text at the top notes the total number of particles in each volume.

Because of this, and because we have demonstrated in other sections that SPHENIX is able to converge on shock problems at faster than first order, we believe the slow convergence on the Sedov problem to be of little importance in practical applications of the scheme.

APPENDIX E: CONDUCTION IN THE NOH PROBLEM

In Section 5.4, we presented the Noh problem and showed that the SPHENIX scheme (like other SPH schemes in general) struggles to capture the high density in the central region due to so-called 'wall heating'.

The SPHENIX scheme includes a switch to reduce artificial conduction in viscous flows. This is, as presented in Section 4, to allow for the capturing of energetic feedback events. It does, however, lead to a minor downside; the stabilizing effect of the conduction in these shocks is almost completely removed. Usually, the artificial conduction lowers the dispersion in local internal energy values, and hence pressures, allowing for a more regular particle distribution.

In Fig. E1, we show three re-simulations of the Noh problem (at 256^3 resolution) with three separate schemes. The first, the full SPHENIX scheme, is simply a lower resolution version of Fig. 10. The second, 'No Conduction Limiter', is the SPHENIX scheme, but with equation (34) removed; i.e. the particle-carried artificial conduction coefficient depends solely on the local internal energy field (through $\nabla^2 u$ and equation (28)), instead of also being mediated by the velocity divergence field. The final case, 'Fixed $\alpha_D = 1.0$', shows the case where we remove all conduction switches and use a fixed value for the conduction α_D of 1.0. The former two look broadly similar, suggesting that the post-shock region is not significantly affected by the additional SPHENIX conduction limiter.

The final panel, however, shows the benefits available to a hypothetical scheme that can remove the artificial conduction switch; the central region is able to hold a significantly higher density thanks to energy being conducted out of this region, allowing the pressure to regularize. In addition to the above, this case shows significantly weaker spurious density features (recall that the post-shock, high-density, region should have a uniform density) because these have been regularized by the conduction scheme.

We present this both to show the drawbacks of the SPHENIX artificial conduction scheme, and to show the importance of demonstrating test problems with the same switches that would be used in a production simulation.

APPENDIX F: BLOB TEST

In Fig. 15, we demonstrated the performance of the SPHENIX scheme on an example 'blob' test. Here, we show how the same initial conditions are evolved using two schemes: a 'traditional SPH' scheme with fixed artificial viscosity ($\alpha_V = 0.8$) and no artificial conduction (e.g. Monaghan 1992),[9] and a SPH-ALE (Vila 1999) scheme similar to GIZMO-MFM[10] (Hopkins 2015) with a diffusive slope limiter. This is in an effort to demonstrate how the initial conditions are evolved with a minimally viable non-diffusive scheme, through to what could be considered the most diffusive viable scheme.

Fig. F1 shows the result of the blob test with the traditional SPH scheme. Here, as expected, there is a severe lack of mixing, with the artificial surface tension holding the blob together even at the highest resolutions. The lack of phase mixing also contributes to a lack of overall mixing, with the stripped trails (shown most clearly at $t = 3\tau_{\rm KH}$) adiabatically expanding but crucially remaining distinct from the hot background medium.

Fig. F2 shows the result of the blob test with the SPH-ALE (GIZMO-MFM) scheme. This scheme is known to be highly diffusive (due to the less conservative slope limiter employed in the SWIFT implementation). This follows closely the results seen in e.g. Agertz et al. (2007) for diffusive grid-based codes. Here, the blob is rapidly shattered, and then dissolves quickly into the surrounding media, especially at the lowest resolutions.

The SPHENIX results in Fig. 15 showed that the blob mixed with the surrounding media, but at a less rapid rate than in the SPH-ALE case. This is somewhat expected, given the trade-off required in the artificial conduction switches (equation 34). We do note, however, that no analytical solution exists for the blob test, and as such all of these comparisons may only be made qualitatively.

In Fig. F3, we examine the effect of removing the conduction limiter from the SPHENIX implementation (i.e. equation 34 is removed, allowing α_D to vary irrespective of the values of α_V). We see that the inclusion of the limiter does slightly reduce the rate of initial mixing within the blob, but that the effect of the limiter is not particularly strong within this case.

[9]The `minimal` scheme in SWIFT.
[10]The `gizmo-mfm` scheme in SWIFT with a HLLC Riemann solver.

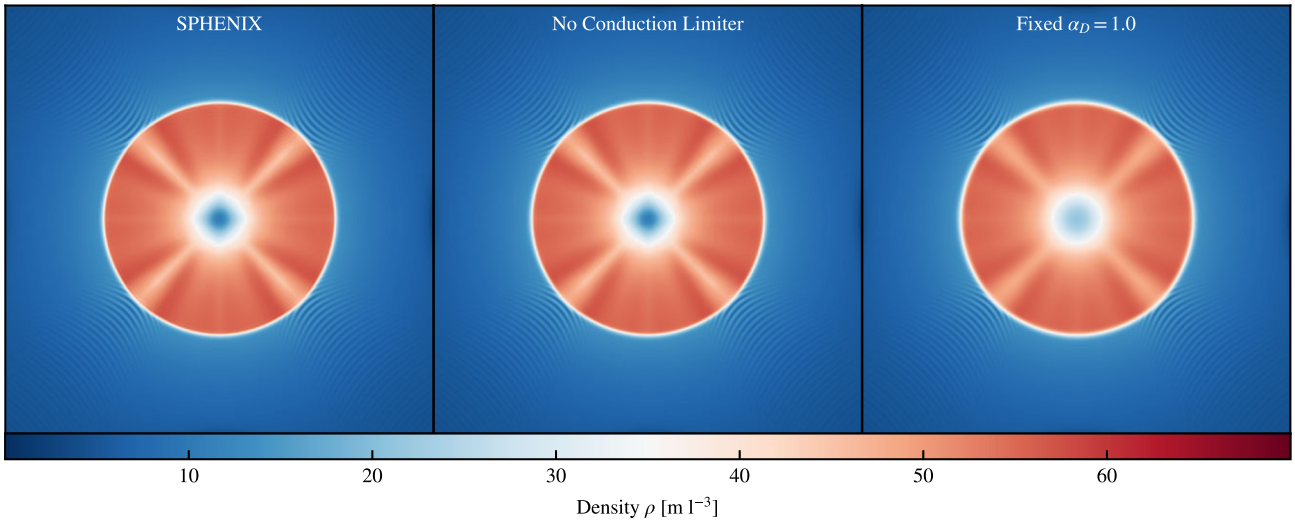

Figure E1. Density slice through the centre of the Noh problem (analogue of Fig. 10) shown for three different artificial conduction schemes (see text). The colour bar is shared between all, and they all use the same, 256^3, initial condition, and are also all shown at $t = 0.6$. The case with the fixed, high, conduction coefficient (right) shows the smallest deviation in density in the centre, as the conduction can treat the wall heating present in this test.

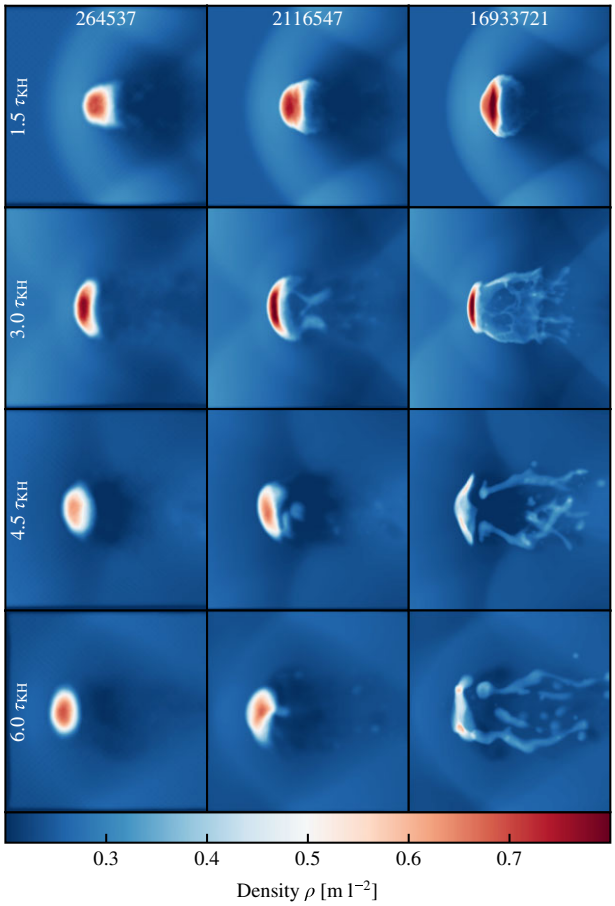

Figure F1. A repeat of Fig. 15 but using a 'traditional' SPH scheme without diffusive switches.

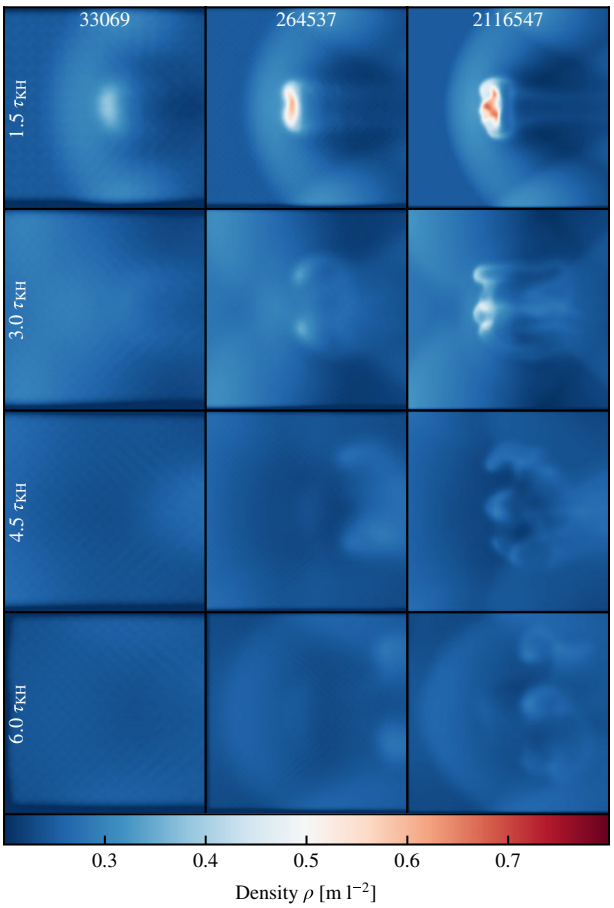

Figure F2. A repeat of Fig. 15 but using an SPH-ALE scheme with a diffusive slope limiter. Note, however, that this is one step lower in resolution, due to the additional computational cost required to perform a simulation including a Riemann solver.

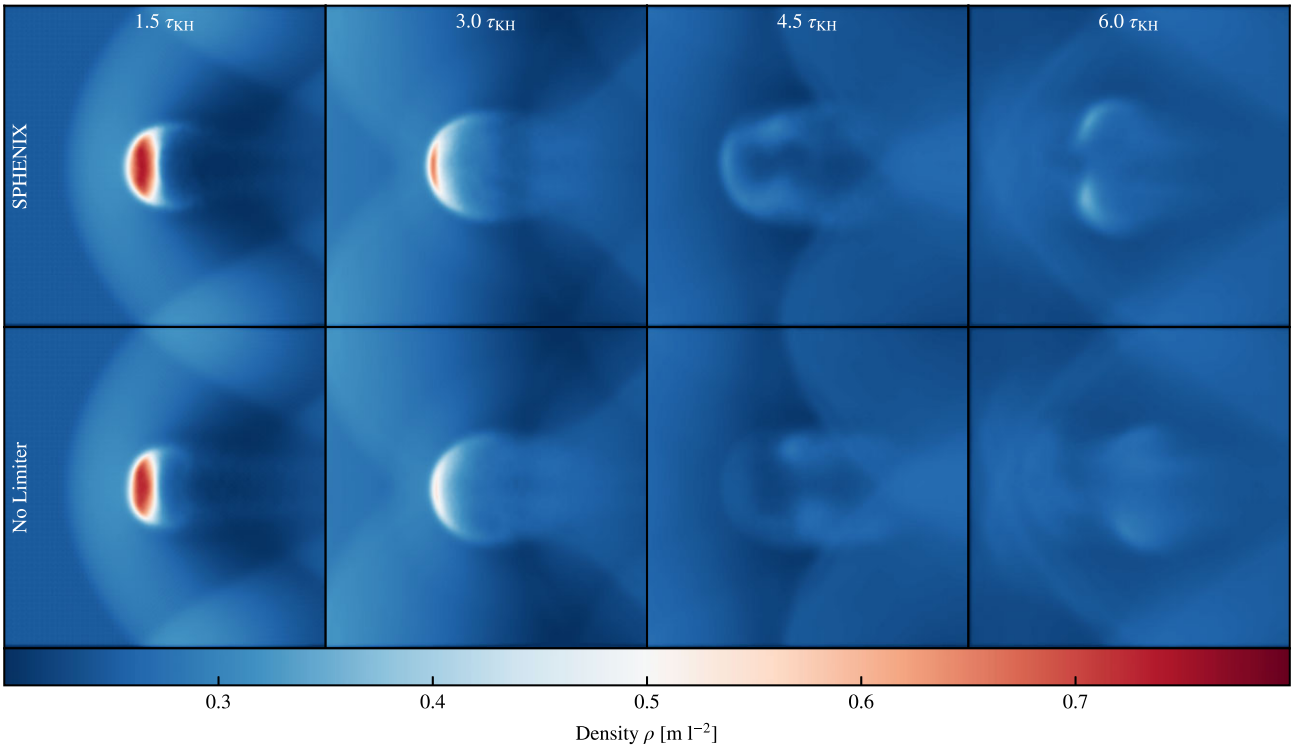

Figure F3. The evolution of a single blob (using the medium resolution, 2116 547 particle, initial conditions from Fig. 15), to illustrate the effect of turning off the conduction limiter (equation 34; bottom row) in comparison to the full SPHENIX scheme (top row). The limiter suppresses some of the initial mixing during the cloud crushing, but does not cause significant qualitative changes in the mixing of the cloud.

This paper has been typeset from a TEX/LATEX file prepared by the author.