

Choosing the best flavour of SPH for astrophysics problems: more complex may not be better

Josh Borrow

Institute for Computational Cosmology
Durham University
Durham, DH1 3LE, UK
joshua.borrow@durham.ac.uk

Bert Vandebroucke

SUPA, School of Physics & Astronomy
University of St. Andrews
St. Andrews, KY16 9SS, UK

Matthieu Schaller

Leiden Observatory
Leiden University
NL-2300 RA Leiden, NL

Abstract—We present SPH(-ALE) schemes that are used in cosmological simulations, and use the Sedov-Taylor blastwave, Evrard collapse, Gresho-Chan vortex, and Sod shock tube at relevant resolutions to determine if the added computational cost of the use of Riemann solvers in cosmological applications is warranted. We use the SWIFT code to enable fair comparisons, using the same neighbour finding framework, data structures, and kernel implementations, for all schemes to enable fair comparisons. We crucially do not individually tune parameters for these tests. We find that, with the exception of strong shocks, at a fixed *runtime*, modern SPH schemes with variable artificial viscosity and thermal diffusion are still more than viable for the adaptive regimes within which cosmology simulations operate.

I. INTRODUCTION

Cosmological simulations have been used in astronomy for decades to provide predictions for observables based on underlying theoretical models. With these simulations, astronomers can predict the abundances, sizes, and formation times for many classes of astronomical objects, from small populations of stars all the way up to large groups and clusters of galaxies. The first cosmological simulations originated in the late 80's and early 90's [1] and only solved for the gravitational dynamics. This is a reasonable assumption, with 80% of the mass of the Universe made up of 'dark matter' that behaves as a collisionless fluid [2], and has the advantage of avoiding 'messy' gas physics. On top of this matter density field, a 'semi-analytic' model was used to study the formation and evolution of the visible matter component ([3] and references). Hydrodynamics introduced to these gravity solvers by the TreeSPH code [4], however at this time computational resources were not developed enough to run a full hydrodynamical model of a representative portion of the Universe. A revolution in the field took place in the early 10's, with Horizon-AGN, Illustris, and EAGLE [5]–[7] running cosmological hydrodynamics simulations that were able to accurately capture the huge dynamic range of 10^{11} in density required to resolve both the large-scale structure of the Universe and the formation of individual galaxies simultaneously [8]. These suites also include extensive 'sub-grid' (below the resolution scale) models for galaxy formation, which must be developed, or at the very least tuned, for a specific hydrodynamics solver. In these simulations that use

around $1000^3 - 2000^3$ resolution elements for the whole computational domain, bound structures (galaxies) are resolved by around 10^{4-5} particles.

In cosmology, there has been a push in recent years to adopt more complex hydrodynamical models. Early simulations used basic Density-Entropy SPH (see [9]) or grid-based techniques (e.g. the second-order Godunov solver in [10]). The focus in recent years has shifted to moving-mesh methods (with a large community using the AREPO code [11]), and SPH-ALE methods with a finite mass solver (typically referred to as GIZMO-MFM in the astronomy community, see [12]). It is still unclear, however, that these methods present any tangible improvement in hydrodynamical accuracy, considering both their increased computational cost (which could be spent integrating more resolution elements) and the low particle count with which the actual structures that are studied in the simulation are typically resolved.

In higher resolution simulations, such as those in the FIRE-2 project [13], which follow 10^{7-8} particles for a single bound structure, SPH-ALE schemes are usually used. Lower-resolution simulations such as those in the EAGLE project [14], where a similar structure is resolved by 10^{4-5} particles, typically use modern SPH schemes that do not include a Riemann solver. There has until now not been a systematic study of which scheme is most profitable in which regime, and there has been no effort to make a specific, sensible, choice. Previously, different schemes were implemented in different codes, with vastly different neighbour-search algorithms and data structures, making comparisons problematic. In this work, we use the SWIFT code [8], [15], with the same neighbour finding framework, data structures, and kernel implementation, for all schemes. We also ensure that the parameters in the hydrodynamics schemes remain fixed for all tests, to ensure that the results are representative of our production simulations.

II. RELAXED INITIAL CONDITIONS

The hydrodynamics tests in this paper are all presented in three dimensions, as opposed to the typical one dimension, to enable fairer comparisons with the cosmological simulations against which we wish to calibrate. Doing tests in 3D makes the process of generating initial conditions significantly more

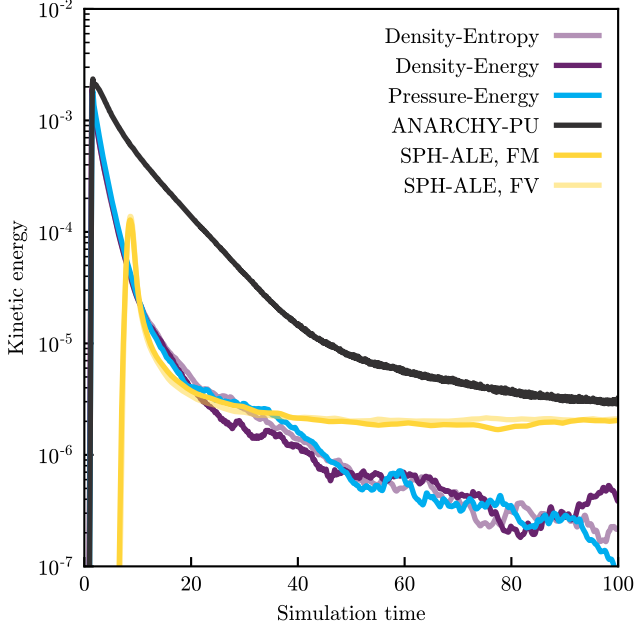


Fig. 1. Total kinetic energy of the particles in the glass-making script as they transition towards a “glass”-like structure. All simulations are stopped at $t = 100$, and these are taken to be settled states of the particles ready for use in the below hydrodynamics tests. This image is shown for the 32^3 particle resolution with the Wendland-C2 kernel.

complex, as to prevent the injection of spurious kinetic (and thermal) energy, they must first be relaxed.

To create these initial condition files, a “glass” file is used. To generate these glass files with SWIFT, we set up an initial grid field, evenly spacing the positions of each particle within the region $[0, 0, 0]$ to $[1, 1, 1]$ (cm). This cube is then evolved with the code using periodic boundary conditions for 100 s of internal code time, long enough for the kinetic energy of the particles in the volume to settle to a stable state (see Fig. 1).

This procedure must be repeated for all glass resolutions, and for all hydrodynamics schemes (with a particular emphasis on different kernel choices), as these all have different equilibrium particle configurations.

III. SCHEMES

This paper compares several schemes: Density-Energy SPH (DU), Density-Entropy SPH (DA), Pressure-Energy SPH (PU), Pressure-Energy SPH with modern switches for both artificial conduction and viscosity (ANARCHY-PU), finite mass SPH-ALE (GIZMO-MFM), and finite volume SPH-ALE (GIZMO-MFV). We pair these with three kernels, a cubic spline, quintic spline, and the Wendland-C2 kernel, as described in [16]. Throughout the paper we take a ratio of smoothing length to mean inter-particle separation $\eta = 1.238$ to give 48 neighbours with the cubic spline for compatibility with the well known Gadget-2 cosmological simulation code, for which this parameter and kernel are the default. For all schemes, we use multiple time-stepping and variable smoothing lengths, as would be used in a production simulation.

A. Basic SPH schemes

1) *Choice of artificial viscosity*: For the three schemes described in this sub-section, DU, DA, and PU, we use the

same, fixed, artificial viscosity prescription as described in [17], along with a Balsara switch [18].

This uses a viscous equation of motion for all particles,

$$\frac{d\mathbf{v}_i}{dt} = - \sum_j \frac{1}{2} \nu_{ij} (\nabla W(r_{ij}, h_i) + \nabla W(r_{ji}, h_j)), \quad (1)$$

which is a sum over neighbouring particles j , with h the smoothing length, m the particle mass, r_{ij} the separation between the particles, and W the SPH kernel. The dimensionless viscosity parameter $\alpha = 0.8$. The viscosity coefficient is defined as follows for each pair of particles

$$\nu_{ij} = - \frac{1}{2} \frac{\alpha (B_i + B_j) \mu_{ij} v_{\text{sig},ij}}{\rho_i + \rho_j}. \quad (2)$$

The signal velocity for each particle is given by the maximum of $v_{\text{sig},ij} = c_i + c_j - 3\mu_{ij}$ over its neighbours with

$$\mu_{ij} = \begin{cases} \frac{\mathbf{v}_{ij} \cdot \mathbf{x}_{ij}}{|\mathbf{x}_{ij}|} & \mathbf{v}_{ij} \cdot \mathbf{x}_{ij} < 0, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

The Balsara switch [18] is implemented as follows,

$$B_i = \frac{|\nabla \cdot \mathbf{v}_i|}{|\nabla \cdot \mathbf{v}_i| + |\nabla \times \mathbf{v}_i| - 10^{-4} c_i / h_i} \quad (4)$$

with the spatial derivatives calculated in the usual way [19].

Along with this equation of motion, there is an associated equation of motion for the particle-carried internal energy u (or entropy A), omitted here for brevity.

2) *Density-Energy SPH (DU)*: This implementation of SPH uses the smooth density for each particle i ,

$$\rho_i = \sum_j m_j W(r_{ij}, h_i). \quad (5)$$

This density is then used with an equation of state, $P_i = (\gamma - 1)u_i\rho_i$, along with a particle-carried internal energy per unit mass u_i , to produce the following SPH equation of motion:

$$\frac{d\mathbf{v}_i}{dt} = - \sum_j m_j \left[f_i \frac{P_i}{\rho_i^2} \nabla W(r_{ij}, h_i) + f_j \frac{P_j}{\rho_j^2} \nabla W(r_{ji}, h_j) \right], \quad (6)$$

with the f factors

$$f_i = \left(1 + \frac{h_i}{n_d \rho_i} \frac{\partial \rho_i}{\partial h_i} \right)^{-1}, \quad (7)$$

with n_d the number of spatial dimensions, which is introduced to account for the variable smoothing lengths required for adaptive astrophysical problems.

3) *Density-Entropy SPH (DA)*: This implementation of SPH is very similar to the Density-Energy implementation described above, but tracks entropy $A_i = (\gamma - 1)u_i/\rho^{\gamma-1}$ as the thermodynamic variable, rather than internal energy. This scheme requires no equation of motion for entropy in inviscid flows as this is conserved under adiabatic expansion. The choice of A as the thermodynamic variable is very popular in astronomy due to its use in the Gadget-2 code [9].

4) *Pressure-Energy SPH (PU)*: Pressure-Energy SPH differs from the above Density-Energy SPH in that it tracks an additional smoothed quantity, the smooth pressure [20],

$$\bar{P}_i = \sum_j (\gamma - 1) u_j m_j W(r_{ij}, h_i), \quad (8)$$

that enters the equation of motion

$$\frac{d\mathbf{v}_i}{dt} = - \sum_j (\gamma - 1)^2 m_j u_j u_i \left[\frac{f_{ij}}{\bar{P}_i} \nabla W(r_{ji}, h_j) + \frac{f_{ji}}{\bar{P}_j} \nabla W(r_{ji}, h_j) \right]. \quad (9)$$

The factors entering the equation of motion due to the variable smoothing lengths in this case are more complex, with

$$f_{ij} = 1 - \left[\frac{h_i}{n_d(\gamma - 1)\bar{n}_i m_j u_j} \frac{\partial \bar{P}_i}{\partial h_i} \right] \left(1 + \frac{h_i}{n_d \bar{n}_i} \frac{\partial \bar{n}_i}{\partial h_i} \right)^{-1}. \quad (10)$$

where \bar{n} is the local number density of particles. These factors also enter the equation of motion for u .

The sound-speed in Pressure-Energy requires some consideration. To see what the ‘correct’ sound-speed is, consider the equation of motion (Eq. 9) in contrast to the EoM for Density-Energy SPH (Eq. 1). For Density-Energy SPH,

$$\frac{d\mathbf{v}_i}{dt} \sim \frac{c_{s,i}}{\rho_i} \nabla_i W_{ij},$$

and for Pressure-Energy SPH,

$$\frac{d\mathbf{v}_i}{dt} \sim (\gamma - 1)^2 \frac{u_i u_j}{\bar{P}_i} \nabla_i W_{ij}.$$

From this it is reasonable to assume that the sound-speed, i.e. the speed at which information propagates in the system through pressure waves, is given by the expression

$$c_{s,i} = (\gamma - 1) u_i \sqrt{\gamma \frac{\rho_i}{\bar{P}_i}}. \quad (11)$$

This expression is dimensionally consistent with a sound-speed, and includes the variable used to evolve the system, the smoothed pressure \bar{P} . However, such a sound-speed leads to a considerably *higher* time-step in front of a shock wave (where the smoothed pressure is higher, but the SPH density is relatively constant), leading to integration problems. An alternative to this is to use the smoothed pressure in the place of the ‘real’ pressure. Whilst it is well understood that \bar{P} should not be used to replace the real pressure in general, here (in the sound-speed) it is only used as part of the time-stepping condition. Using the ‘smoothed’ sound-speed

$$\bar{c}_{s,i} = \sqrt{\gamma \frac{\bar{P}_i}{\rho_i}} \quad (12)$$

instead of Eq. 11 leads to a much improved time-stepping condition that actually allows particles to be woken up before being hit by a shock.

B. ANARCHY-PU SPH

The ANARCHY scheme, known for its use in the EAGLE simulations [7], [14], [21], builds on top of the Pressure-SPH discussed above. It also includes a simplified version of the artificial viscosity scheme from [22], and a simple diffusion term [23]. The scheme implemented in SWIFT differs from the one used in the original EAGLE simulations, using the internal energy as the thermodynamic variable as opposed to the particle entropy.

The artificial viscosity scheme is similar to the above, but now the coefficient α varies with time. The aim here is to have the artificial viscosity peak at a shock front, and then decay afterwards to reduce spurious dissipation in non-shocking regions. To that end, we use a pre-shock indicator

$$S_i = -h_i^2 \min(\dot{\nabla} \cdot \mathbf{v}_i, 0) \quad (13)$$

for each particle. This, the time differential of the velocity divergence, is calculated implicitly from the value at the previous step, i.e.

$$\dot{\nabla} \cdot \mathbf{v}_i(t) = \frac{\nabla \cdot \mathbf{v}_i(t) - \nabla \cdot \mathbf{v}_i(t - \Delta t)}{\Delta t}. \quad (14)$$

The velocity divergence is again calculated in the regular SPH way, with no ‘special’ or ‘improved’ estimate used. The viscosity for each particle α_i is set using the following logic:

$$\alpha_i \rightarrow \begin{cases} \alpha_{\text{loc},i} & \text{if } \alpha_i < \alpha_{\text{loc},i}, \\ \alpha_{\text{loc},i} + (\alpha_i - \alpha_{\text{loc},i}) e^{-dt/\tau_i} & \text{if } \alpha_i > \alpha_{\text{loc},i}, \\ \alpha_{\text{min}} & \text{if } \alpha_i < \alpha_{\text{min}}, \end{cases} \quad (15)$$

with $\alpha_{\text{loc},i} = \alpha_{\text{max}} S / (S + v_{\text{sig},i}^2)$ and $\tau_i = \ell c_{s,i} / H_i$, where $\alpha_{\text{max}} = 2.0$, $\alpha_{\text{min}} = 0.0$, $\ell = 0.25$, are dimensionless parameters, and H_i the kernel support radius. Alongside this, we again use the Balsara switch to reduce dissipation in shear flows.

The final ingredient, which was added to ensure that clusters have flat entropy profiles in adiabatic cosmological simulations [24], and to ensure that Kelvin-Helmholtz instabilities have a smooth profile, is thermal diffusion. This thermal diffusion has a switch that aims to ensure that it is only activated in cases of extreme energy gradients, with

$$\frac{d\alpha_{D,i}}{dt} = \beta_D H_i \frac{\nabla^2 u_i}{\sqrt{u_i}} + \frac{\alpha_{D,i} - \alpha_{D,\text{min}}}{\tau_{D,i}}, \quad (16)$$

with $\alpha_{D,\text{min}} = 0$, $\beta_D = 0.01$, and $\tau_{D,i} = v_{\text{sig},i} / H_i$ included to allow for decay away from discontinuities. This is then bounded by $\alpha_{D,\text{max}} = 1$, however typical values are of order 0.01, which is much lower than the typical value in other codes. The actual diffusion is then applied to the particles with the following equation of motion:

$$\frac{du_i}{dt} = \sum_j \frac{\alpha_{D,i} + \alpha_{D,j}}{\rho_i + \rho_j} v_{D,ij} (u_i - u_j) \bar{W}_{ij} \quad (17)$$

with $\bar{W}_{ij} = \frac{1}{2} (\nabla W(r_{ij}, h_i) + \nabla W(r_{ji}, h_j))$ and the diffusion velocity $v_{D,ij} = \max(c_i + c_j + \mu_{ij}, 0)$.

C. SPH-ALE Schemes

Both SPH-ALE flavours evolve a set of conservation laws for the mass, momentum and energy of the particles, as detailed in e.g. [11], [12], [25], [26]. Conserved variables are converted into primitive variables (density, velocity, internal energy, pressure) using an SPH volume estimate:

$$V_i = \frac{1}{\sum_j W(r_{ij}, h_i)}. \quad (18)$$

The change in conserved variables is then computed using a generalised version of the finite volume equation in which conserved variables are *exchanged* between neighbouring particles through an interface with a formal surface area computed from the kernel gradients. The fluxes themselves are based on the solution of a Riemann problem at the abstract interface using an HLLC Riemann solver.

The two flavours of SPH-ALE differ in their assumption about the movement of the abstract interface:

1) *GIZMO-MFM*: assumes the interface moves with the velocity of the contact discontinuity in the solution of the Riemann problem. Since no mass fluxes are possible through the contact discontinuity, this effectively cancels mass flux between particles. Note that this assumption for the interface velocity cannot be enforced in the particle movement; it will to second order accuracy be satisfied if the particles are moved with the local fluid velocity, but could lead to artificial suppression of mass fluxes in regimes of strong shock waves.

2) *GIZMO-MFV*: assumes that the interface velocity moves with the geometrically averaged speed of the neighbours that share the interface. Mass fluxes are allowed, so that the movement of the particles no longer needs to be fixed to the local fluid velocity. Because of the additional flux, this scheme requires more variables to be stored per particle.

IV. HYDRODYNAMICS TESTS

Below we present four hydrodynamics tests, with particle distributions shown at typical resolutions for resolved objects in cosmological simulations. It is important to note that this paper does not claim that these are the ‘best’ results possible, or that these schemes are perfect in any way. We aim to produce results that accurately characterise the solvers that are used in production simulations, noting that a similar resolution to a resolved object in cosmological simulations lies between 32^3 (3×10^4) and 64^3 (3×10^5). We hence choose 32^3 as the fiducial resolution for the tests below.

To show convergence, we use the L1 norm per particle

$$L1 = \frac{\sum_i |Q_{i,true} - Q_{i,sim}|}{n}, \quad (19)$$

with the sum over n particles i and their analytic solution at that point Q_{true} compared against the particle values Q_{sim} . We also choose to show convergence as a function of simulation time-to-solution (runtime), as this also accounts for the higher computational cost of the ALE schemes.

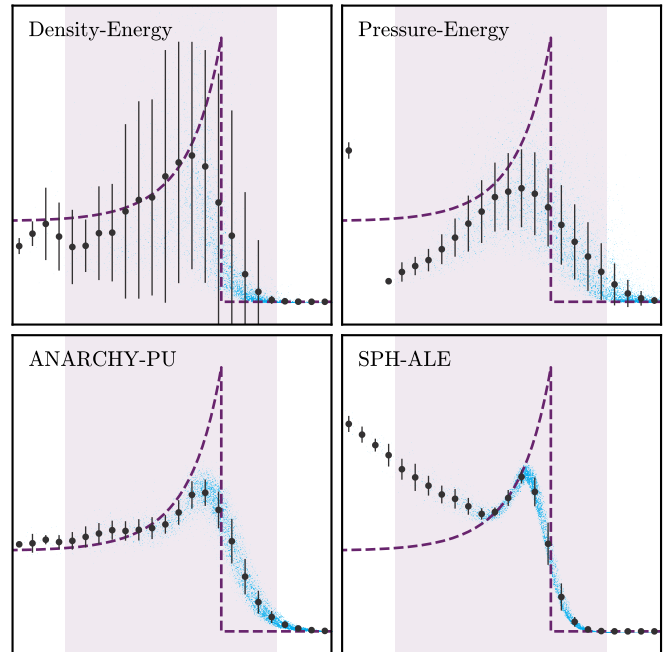


Fig. 2. The pressure as a function of radius in the Sedov-Taylor blastwave for a selection of schemes. The scatter in the background shows all particles for this 32^3 run with the Wendland-C2 kernel, with the over-plotted black points showing the mean in bins along with the 1σ scatter. The purple dashed line shows the exact solution, with the purple shaded region showing where the L1 norm in Fig. 3 is calculated. All are shown at the same time $t = 0.05$.




Fig. 3. L1 convergence for the pressure as shown in Fig. 2 for the Sedov-Taylor blastwave. Greyed out lines and data show the Density-Entropy variant of SPH (purple), and the finite-volume variant of the SPH-ALE scheme (yellow). The individual test simulations are plotted as points (and the fiducial 32^3 shown with stars), with straight lines being added as linear fits to the data. Note that the horizontal axis here is *total runtime* on 28 cores.

A. Sedov-Taylor Blastwave

The Sedov-Taylor blastwave (Sedov blast) [27], [28] 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.

1) *Initial Conditions:* The relevant glass file from §II is chosen, and the particle properties are 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.

2) *Results:* The SPH-ALE scheme outperforms all of the SPH schemes, including the (‘modern’) ANARCHY-PU, in the convergence with run-time test (Fig. 3). Not only does it converge significantly faster, but it even provides a lower norm at poor resolution than the SPH schemes can achieve even at high resolution. Despite this, the schemes perform similarly at our fiducial, 32^3 resolution, with a particular note on the similarity between the SPH-ALE and ANARCHY-PU scheme. Note that schemes have a more compressed runtime range due to varying numbers of steps taken and the time-step hierarchy mechanism described in [4], [8].

The particle distributions tell a marginally different picture; the SPH schemes better capture the energy distribution in the post-shock region. Most convergence studies performed on the Sedov blast use the density distribution of the particles, however in cosmological simulations where Sedov blasts are used to model the feedback processes from supernovae, the energy profile in the post-shock region must be correct. In this region, radiative cooling will take place (as a sub-grid process), for which the cooling rate is dependent on the current energy of the particles. This cooling rate can also be converted to a radiative flux for mock observations of galaxies. For the ALE scheme, the cooling rate in the post-shock region will be estimated to be significantly higher than the true value, leading to incorrect fluxes around these important star-forming regions.

B. Gresho-Chan Vortex

The Gresho-Chan vortex [29] is typically used to test for the conservation of vorticity and angular momentum, and is usually performed in 2D. Here, we extend it to 3D, with a vortex tube.

1) *Initial Conditions:* The initial conditions use the appropriate glass file from §II, and treat the gas with an adiabatic index $\gamma = 5/3$, constant density $\rho_0 = 1$, in a cube of side-length 1. The particles are given azimuthal velocity

$$v_\phi = \begin{cases} 5r & r < 0.2 \\ 2 - 5r & 0.2 \leq r < 0.4 \\ 0 & r \geq 0.4 \end{cases} \quad (20)$$

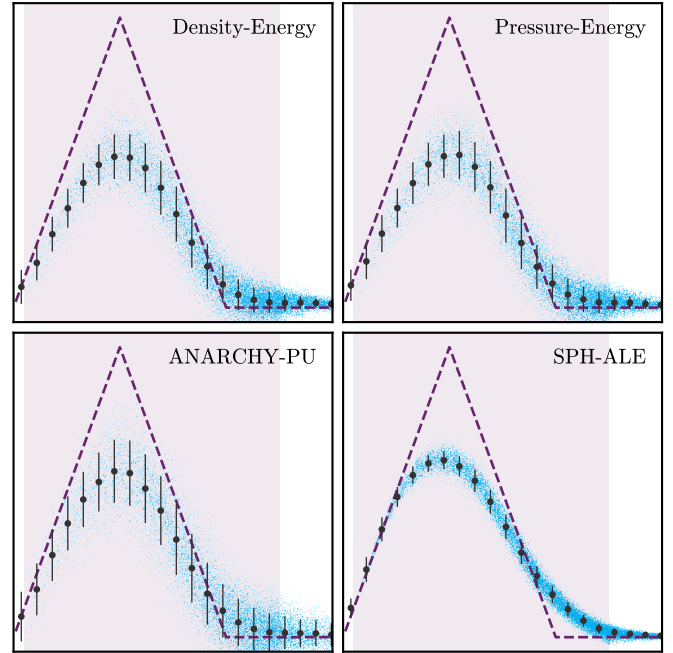


Fig. 4. Azimuthal velocity as a function of radius, with the colours in the figure the same as Fig. 2. The schemes are shown at $t = 1$, which is just before the regular SPH schemes completely break down (at approximately $t = 3$, see [12]). The test is performed in 3D, at the fiducial 32^3 resolution, with the Wendland-C2 kernel.




Fig. 5. The L1 convergence for the azimuthal velocity in the Gresho vortex test at $t = 1$, with lines coloured the same as in Fig. 3.

with the pressure in equilibrium as

$$P_0 = \begin{cases} 5 + 12.5r^2 & r < 0.2 \\ 9 + 12.5r^2 - 20r + 4 \log(5r) & 0.2 \leq r < 0.4 \\ 3 + 4 \log(2) & r \geq 0.4 \end{cases} \quad (21)$$

where here $r^2 = x^2 + y^2$ from the box centre.

2) *Results:* From the particle distributions for v_ϕ in Fig. 4 it is clear that the SPH schemes do not perform well on this test, as expected. The ANARCHY-PU scheme is able to produce

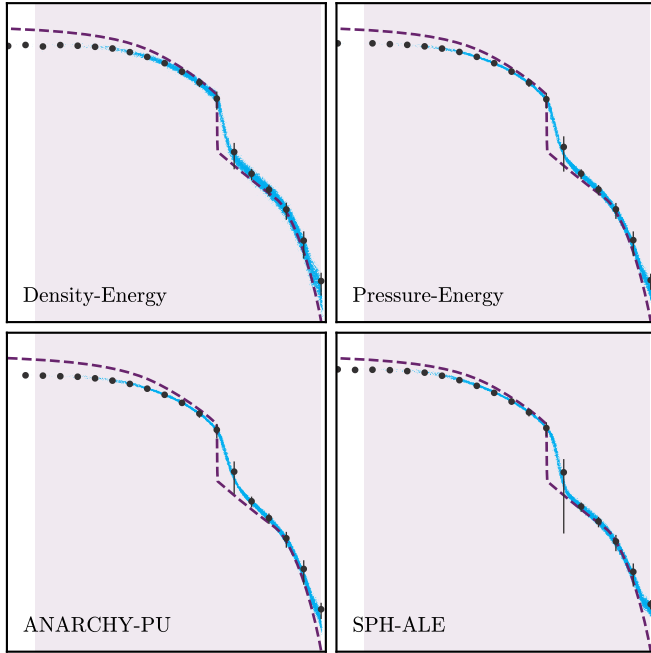


Fig. 6. The pressure profile of the Evrard sphere mid-collapse at $t = 0.8$, with the colours the same as in Fig. 2. Shown here is the result at the fiducial 32^3 resolution with the Wendland-C2 kernel.

a slightly closer mean distribution due to a lower $\alpha \approx 0.1$ in the inner regions, but the Balsara switch does not manage to fully suppress this for the converging flow. The α coefficient here is also very noisy, with a scatter from $0.0 < \alpha < 0.2$, leading to a high level of scatter in the velocity profile, which is then magnified further as the simulation progresses due to α being evolved as function of the velocity divergence.

After the $t = 1$ solution shown, which is approximately one rotation of the vortex at the peak, the SPH solutions collapse completely, with the SPH-ALE result remaining stable for much longer. The SPH-ALE result does well here, but is still severely limited at the velocity peak by the performance of the slope limiter. This is shown in the convergence profile, Fig. 5, where at these low resolutions we see that the ANARCHY-PU SPH converges just as quickly as the SPH-ALE scheme, as well as providing the result in a considerably lower wall-clock time.

C. Evrard Collapse

The Evrard collapse [30] test takes a large sphere of self-gravitating gas, at low energy and density, that 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.

1) *Initial Conditions:* Gas particles are set up with an adiabatic index of $\gamma = 5/3$, mass $M = 1$, 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 there are no effects from the periodic boundary. Unfortunately, due to the non-uniform density profile, it is considerably more challenging to provide relaxed initial conditions (or use a

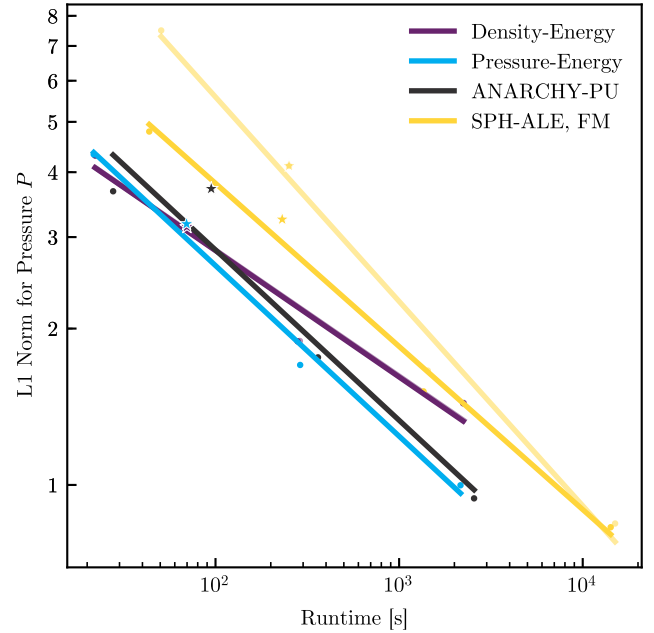


Fig. 7. The L1 convergence for the pressure of the particles for the Evrard collapse at $t = 0.8$, with lines coloured the same as in Fig. 3. Note how similar all of the convergence profiles are. In this calculation, the gravity solver dominates the runtime, allowing the SPH-ALE schemes to gain some ground back against the schemes without Riemann solvers.

glass file). Here, positions are simply drawn randomly to produce the required density profile, with the same set of initial conditions used for all schemes.

2) *Results:* After a short period of in-fall, a shock is set up that is predictably smoothed within all of the particle-based schemes, as shown in Fig. 6. All schemes provide a roughly similar picture here, with pressure chosen as the comparison variable to ensure that there is no spurious transfer of energy from the central regions to the outskirts for the ANARCHY scheme that includes thermal diffusion. The schemes all capture the accretion shock similarly, with the ANARCHY-PU scheme resulting in the smoothest shock, and the Density-Energy scheme resulting in the sharpest. This is likely due to spurious artificial viscosity activation in the outer ($r > 0.5$) region, leading to $\alpha > 1.0$ and excess diffusivity.

Of particular interest in this case is the convergence with runtime, shown in Fig. 7, with all lines overlaid on top of each other. The marginally higher computational cost per step (and larger number of steps) ensures that the SPH-ALE has the highest runtime normalisation, even despite the added fixed cost of the gravity solver.

This test shows the limitations of an SPH-ALE scheme at the very low resolutions that we target. The convergence here is still very much limited by the width of the shock front, which in any scheme can only be resolved to within a kernel width, irrespective of whether or not a Riemann solver is employed.

D. Sod Shock Tube

1) *Initial Conditions:* The initial conditions for the Sod shock tube use two glass files, one with a step higher resolution (e.g. the 32^3 is paired with the 64^3). The gas is given an

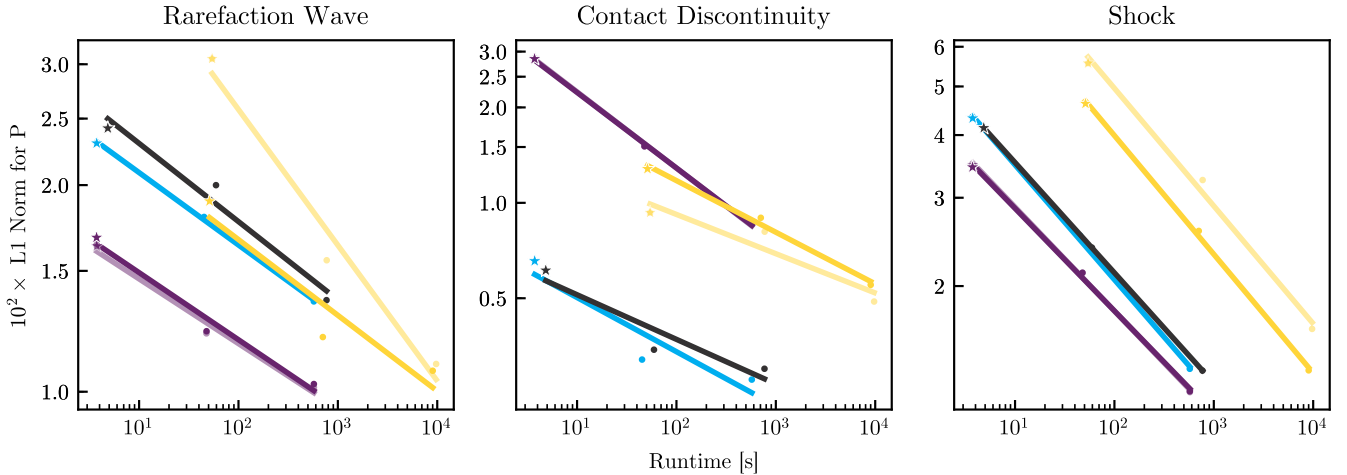


Fig. 8. Convergence for three regions within the Sod Shock tube as a function of total runtime. The three sections have L1 calculated as the mean per particle for the pressure within a fixed 0.1 distance aperture around the peak of the rarefaction wave, the contact discontinuity, and the shock front. Shown here are the results for the Wendland-C2 kernel. Individual lines are coloured similarly to Fig. 3, with the legend omitted here for brevity. Note that the contact discontinuity results for the ALE scheme are due to the use of SPH initial conditions, and do not accurately represent the ability of the scheme to capture such a discontinuity.

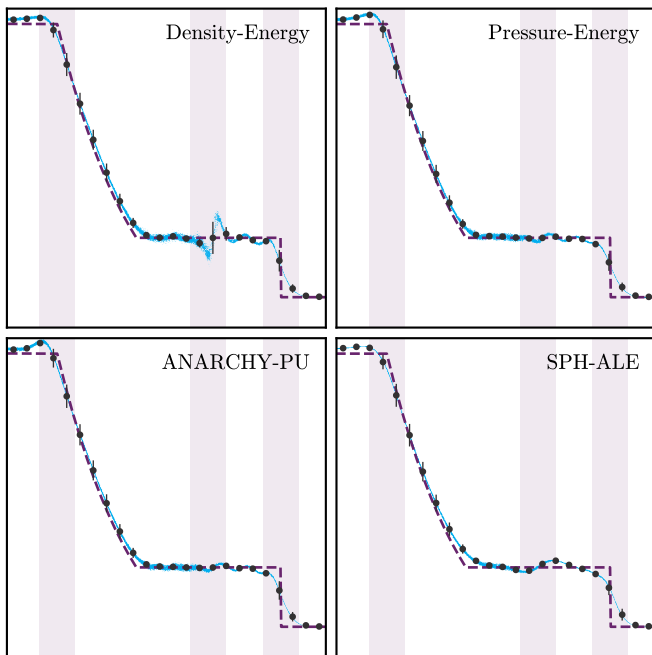


Fig. 9. Images of the pressure distribution of particles within the Sod shock tube at the fiducial 32^3 (plus the 16^3 low-resolution glass, see the text for details) resolution with the Wendland-C2 kernel. The contact discontinuity is present in between the rarefaction wave and shock front, where there is a blip in the Density-Energy particle distribution.

adiabatic index $\gamma = 5/3$, and the two glass files are set up next to each other to form a discontinuity, with the high resolution on the left and the low resolution on the right, with density $\rho_L = 1$, $\rho_R = 1/8$, velocity $v_L = v_R = 0$, and pressure $P_L = 1$, $P_R = 0.1$.

2) *Results:* The pressure distribution in the shock tube, shown in Fig. 9, shows that the overall match between the particle schemes and the reference solution is good. The Density-Energy scheme shows the expected ‘blip’ at the contact discontinuity caused by the large discontinuity in both density and energy here, combined with the fact that pressure

is a derived quantity from these through $P = (\gamma - 1)u\rho$ [23]. The Pressure-Energy based schemes remove this blip by providing a truly smoothed, continuous pressure (Eqn. 8) in this region. The SPH-ALE scheme, at this resolution, provides rather poor results, with the slope limiter having non-negligible effects at the rarefaction wave causing an over-shoot of the pressure. At the contact discontinuity, there is a blip that is due to the setting of mass in the initial conditions as the initial variable, instead of density. This was chosen to enable the same initial conditions for all tests.

As with the Evrard collapse, we see that at this low resolution that all schemes handle the (relatively weak) shock equally well, and that the SPH-ALE scheme cannot provide an improvement over the minimal resolution provided by the kernel width.

This is highlighted in the convergence plots (Fig 8), with the right panel showing equal norms as a function of particle number, with the SPH-ALE schemes moved to the right due to their systematically higher computational cost.

The overshoot in pressure in the rarefaction wave that the SPH-ALE scheme is associated with due to imperfections in the slope-limiting technique causes a systematically higher norm than expected, with even poorer results at the low-pressure end of the wave.

V. CONCLUSIONS AND SUMMARY

In the above, we have presented four standard hydrodynamics tests in an effort to provide recommendations for the astrophysics community. As noted in the introduction, the simulations that are performed by astrophysicists are in a significantly less resolved, more adaptive, and chaotic environment than those in any other field. This means that any scheme that we employ must be robust to outside changes, work well at as low of a resolution as possible, and due to our compute constraints, cost as little as possible.

The above tests show that, in general, a scheme without a Riemann solver is adequate for most cosmological test

cases (i.e. at the resolutions that are studied). Strong shocks are able to benefit from the Riemann solver, with the negative consequence of an incorrect pressure distribution in the post-shock region (with the solvers that are used *today* in cosmological simulations). Future developments in SWIFT, perhaps by using a more sophisticated slope limiter, may enable us to get significantly more benefit out of the SPH-ALE schemes. However, these results cast doubt on the constant push for more complex hydrodynamics solvers in cosmology at commonly used resolutions, and justifies the use of SPH solvers for very large, high particle count, volumes that are only possible with schemes that do not require a Riemann solver due to compute and memory constraints.

VI. ACKNOWLEDGEMENTS

JB is supported by STFC studentship ST/R504725/1. VB is supported by STFC grant ST/M001296/1. MS is supported by MS is supported by VENI grant 639.041.749. This work used the ARCHER UK National Supercomputing Service. 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.

A. Software Citations

This paper made use of the following software packages:

- SWIFT [15]
- python [31], with the following libraries
 - numpy [32]
 - scipy [33]
 - matplotlib [34]
 - py-sphviewer [35]
 - swiftsimio

REFERENCES

- [1] C. S. Frenk, S. D. M. White, M. Davis, and G. Efstathiou, “The Formation of Dark Halos in a Universe Dominated by Cold Dark Matter,” *APJ*, vol. 327, p. 507, Apr 1988.
- [2] Planck Collaboration, “Planck 2015 results. XIII. Cosmological parameters,” *AAP*, vol. 594, p. A13, Sep 2016.
- [3] C. G. Lacey *et al.*, “A unified multiwavelength model of galaxy formation,” *MNRAS*, vol. 462, no. 4, pp. 3854–3911, Nov 2016.
- [4] L. Hernquist and N. Katz, “TRESSPH - A unification of SPH with the hierarchical tree method,” *ApJS*, vol. 70, pp. 419–446, Jun. 1989.
- [5] Y. Dubois *et al.*, “Dancing in the dark: galactic properties trace spin swings along the cosmic web,” *MNRAS*, vol. 444, no. 2, pp. 1453–1468, Oct 2014.
- [6] M. Vogelsberger *et al.*, “Introducing the Illustris Project: simulating the coevolution of dark and visible matter in the Universe,” *MNRAS*, vol. 444, no. 2, pp. 1518–1547, Oct 2014.
- [7] J. Schaye *et al.*, “The EAGLE project: Simulating the evolution and assembly of galaxies and their environments,” *MNRAS*, vol. 446, no. 1, pp. 521–554, 2015.
- [8] J. Borrow, R. G. Bower, P. W. Draper, P. Gonnet, and M. Schaller, “SWIFT: Maintaining weak-scalability with a dynamic range of 10^4 in time-step size to harness extreme adaptivity,” *Proceedings of the 13th SPHERIC International Workshop, Galway, Ireland, June 26-28 2018*, pp. 44–51, Jul 2018.
- [9] V. Springel, “The cosmological simulation code GADGET-2,” *MNRAS*, vol. 364, pp. 1105–1134, Dec. 2005.
- [10] R. Teyssier, “Cosmological hydrodynamics with adaptive mesh refinement. A new high resolution code called RAMSES,” *A&A*, vol. 385, pp. 337–364, Apr. 2002.
- [11] V. Springel, “E pur si muove: Galilean-invariant cosmological hydrodynamical simulations on a moving mesh,” *MNRAS*, vol. 401, no. 2, pp. 791–851, Jan 2010.
- [12] P. F. Hopkins, “A new class of accurate, mesh-free hydrodynamic simulation methods,” *MNRAS*, vol. 450, pp. 53–110, Jun. 2015.
- [13] P. F. Hopkins *et al.*, “FIRE-2 simulations: physics versus numerics in galaxy formation,” *MNRAS*, vol. 480, no. 1, pp. 800–863, Oct 2018.
- [14] M. Schaller, C. Dalla Vecchia, J. Schaye, R. G. Bower, T. Theuns, R. A. Crain, M. Furlong, and I. G. McCarthy, “The EAGLE simulations of galaxy formation: the importance of the hydrodynamics scheme,” *MNRAS*, vol. 454, pp. 2277–2291, Dec. 2015.
- [15] M. Schaller, P. Gonnet, A. B. G. Chalk, and P. W. Draper, “SWIFT: Using task-based parallelism, fully asynchronous communication, and graph partition-based domain decomposition for strong scaling on more than 100,000 cores,” in *Proceedings of the PASC Conference*, Lausanne, Switzerland, Jun. 2016.
- [16] W. Dehnen and H. Aly, “Improving convergence in smoothed particle hydrodynamics simulations without pairing instability,” *MNRAS*, vol. 425, no. 2, pp. 1068–1082, Sep 2012.
- [17] J. J. Monaghan, “Smoothed particle hydrodynamics,” *ARAA*, vol. 30, pp. 543–574, Jan 1992.
- [18] D. S. Balsara, Ph.D. dissertation, -, Jan 1989.
- [19] D. J. Price, “Smoothed particle hydrodynamics and magnetohydrodynamics,” *Journal of Computational Physics*, vol. 231, pp. 759–794, Feb. 2012.
- [20] P. F. Hopkins, “A general class of Lagrangian smoothed particle hydrodynamics methods and implications for fluid mixing problems,” *Monthly Notices of the Royal Astronomical Society*, vol. 428, no. 4, pp. 2840–2856, 2013.
- [21] R. A. Crain *et al.*, “The EAGLE simulations of galaxy formation: calibration of subgrid physics and model variations,” *MNRAS*, vol. 450, pp. 1937–1961, Jun. 2015.
- [22] L. Cullen and W. Dehnen, “Inviscid SPH,” vol. 15, no. June, pp. 1–15, 2010. [Online]. Available: <http://arxiv.org/abs/1006.1524>
- [23] D. J. Price, “Modelling discontinuities and Kelvin Helmholtz instabilities in SPH,” *Journal of Computational Physics*, vol. 227, no. 24, pp. 10 040–10 057, Dec 2008.
- [24] F. Sembolini *et al.*, “nIFTy galaxy cluster simulations - I. Dark matter and non-radiative models,” *MNRAS*, vol. 457, no. 4, pp. 4063–4080, Apr 2016.
- [25] B. B. Moussa, N. Lanson, and J. P. Vila, “Convergence of meshless methods for conservation laws applications to euler equations,” in *Hyperbolic Problems: Theory, Numerics, Applications*, M. Fey and R. Jeltsch, Eds. Basel: Birkhäuser Basel, 1999, pp. 31–40.
- [26] B. Vandenbroucke and S. De Rijcke, “The moving mesh code SHAD-OWFAX,” *Astronomy and Computing*, vol. 16, pp. 109–130, Jul 2016.
- [27] G. Taylor, “The Formation of a Blast Wave by a Very Intense Explosion. I. Theoretical Discussion,” *Proceedings of the Royal Society of London Series A*, vol. 201, no. 1065, pp. 159–174, Mar 1950.
- [28] L. I. Sedov, *Similarity and Dimensional Methods in Mechanics*, 1959.
- [29] P. M. Gresho and S. T. Chan, “On the theory of semi-implicit projection methods for viscous incompressible flow and its implementation via a finite element method that also introduces a nearly consistent mass matrix. II - Implementation,” *International Journal for Numerical Methods in Fluids*, vol. 11, pp. 621–659, Oct 1990.
- [30] A. E. Evrard, “Beyond N-body: 3D cosmological gas dynamics,” *MNRAS*, vol. 235, pp. 911–934, Dec 1988.
- [31] G. van Rossum, “Python tutorial,” Centrum voor Wiskunde en Informatica (CWI), Amsterdam, Tech. Rep. CS-R9526, May 1995.
- [32] T. Oliphant, “NumPy: A guide to NumPy,” USA: Trelgol Publishing, 2006–. [Online]. Available: <http://www.numpy.org/>
- [33] E. Jones, T. Oliphant, P. Peterson *et al.*, “SciPy: Open source scientific tools for Python,” 2001–. [Online]. Available: <http://www.scipy.org/>
- [34] J. D. Hunter, “Matplotlib: A 2d graphics environment,” *Computing In Science & Engineering*, vol. 9, no. 3, pp. 90–95, 2007.
- [35] A. Benitez-Llambay, “py-sphviewer: Py-sphviewer v1.0.0,” Jul. 2015. [Online]. Available: <http://dx.doi.org/10.5281/zenodo.21703>