

Inviscid smoothed particle hydrodynamics

Lee Cullen and Walter Dehnen^{*}

Department of Physics & Astronomy, University of Leicester, Leicester LE1 7RH

Accepted 2010 June 7. Received 2010 June 4; in original form 2009 July 16

ABSTRACT

In smoothed particle hydrodynamics (SPH), artificial viscosity is necessary for the correct treatment of shocks, but often generates unwanted dissipation away from shocks. We present a novel method of controlling the amount of artificial viscosity, which uses the total time derivative of the velocity divergence as shock indicator and aims at completely eliminating viscosity away from shocks. We subject the new scheme to numerous tests and find that the method works at least as well as any previous technique in the strong-shock regime, but becomes virtually inviscid away from shocks, while still maintaining particle order. In particular sound waves or oscillations of gas spheres are hardly damped over many periods.

Key words: hydrodynamics – methods: numerical.

1 INTRODUCTION

Smoothed particle hydrodynamics (SPH) is a Lagrangian method for modelling fluid dynamics, pioneered by Gingold & Monaghan (1977) and Lucy (1977). Instead of discretizing the fluid quantities, such as density, velocity and temperature, on a fixed grid as in Eulerian methods, the fluid is represented by a discrete set of moving particles acting as interpolation points. Because of its Lagrangian nature, SPH models regions of higher density with higher resolution with the ability to simulate large dynamic ranges. This makes it particularly useful in astrophysics, where it is used to model galaxy and star formation, stellar collisions and accretion discs.

The core of SPH is the kernel estimator: the fluid density is *estimated* from the masses m_i and positions \mathbf{x}_i of the particles via¹

$$\hat{\rho}(\mathbf{x}_i) = \sum_j m_j W(|\mathbf{x}_i - \mathbf{x}_j|, h_i), \quad (1)$$

where W is the kernel function and h_i the SPH smoothing length² for the i th particle. Similar estimates $\hat{F}(\mathbf{x})$ for the value of any field $F(\mathbf{x})$ can be obtained from its discretized values F_i . By applying these estimators to the fluid equations governing mass, momentum and energy, discrete equations for the SPH particle positions \mathbf{x}_i and other properties (such as internal energy u_i) can be obtained. Together with an appropriate time integration method, these constitute a concrete SPH scheme.

Unfortunately, this process is not unique and since its inception the SPH method has undergone many refinements such as individual particle smoothing lengths and viscosities, as well as many

alternative derivations of the SPH equations, leading to a plethora of SPH methods. While formally these various schemes differ only in their error terms, their conservation and stability properties can be quite different. This has led to the unfortunate situation that the shortcomings of a few such implementations are often blamed on the general SPH concept per se.

However, Springel & Hernquist (2002) have pointed out that SPH equations derived from a variational principle are not only unique, but also conservative. Such SPH equations are most simply obtained as the Euler–Lagrange equations derived from an SPH Lagrangian \mathcal{L} representing the Lagrangian of the fluid system. Once \mathcal{L} is chosen, the SPH equations follow uniquely (see Appendix A2 for a typical example). Complementing these with a symplectic integrator, such as the standard leapfrog, results in a SPH scheme which by construction conserves the total mass, momentum, angular momentum, energy and entropy.

The conservation of entropy means that SPH is *dissipationless*, as demonstrated in Fig. 1. In real fluids, however, entropy increases in shocks, where particle collisions randomize their velocities generating heat and entropy. This basic collisional mechanism is inherent to all fluids (except for dust and collisionless plasma, which therefore may not be considered fluids) and prevents the flow from becoming multivalued. In SPH *artificial* viscosity is needed to dissipate local velocity differences and convert them into heat, which generates entropy and prevents interpenetration of SPH particles and thus a multivalued flow.

Since the artificial viscosity required for this goal is usually much stronger than the actual physical viscosity, it also causes unphysical dissipation away from shocks. While it may be possible for certain simulations to select the magnitude of the viscosity to minimize such undesired dissipation, in general the adverse effect of artificial viscosity is unknown prior to any simulation and, possibly, even afterwards. For example, when simulating the effect of a perturbing massive body on a pulsating star, it may be very difficult to distinguish this effect from that induced by artificial viscosity. Another

^{*}E-mail: walter.dehnen@astro.le.ac.uk

¹ We use the symbol $\hat{}$ to denote a local *estimate* – in many SPH-related publications the distinction between actual and estimated quantities is not clearly made, confusing the discussion.

² In this study we use the convention that the kernel has finite support of one smoothing length radius, i.e. $W = 0$ for $|\mathbf{x}_i - \mathbf{x}_j| > h$.




Figure 1. A 1D sinusoidal sound wave with velocity amplitude $10^{-4}c$ and $\gamma = 1.4$ propagated for 50 periods with SPH without artificial viscosity using 100 particles and with a grid code (RAMSES; Teyssier 2002) using 128 cells (only every fifth particle or grid cell is plotted). Both methods preserve the wave amplitude and period, demonstrating their dissipationless nature.

example is the case of a differentially rotating disc, where artificial viscosity causes spurious angular momentum transport.

Since viscosity is a dissipative process, the corresponding SPH equations cannot be derived from a variational principle, and we are back to ad hoc methods for deriving them. Most SPH simulations to date still use a rather simple artificial viscosity, which effectively amounts to modelling a viscous fluid and quickly damps away any oscillations, such as sound waves or stellar pulsations, and impedes shear flows. While suggestions have been made to *reduce* such unwanted dissipation, our goal here is to *eliminate* it. To this end we introduce a novel method of controlling the amount of artificial viscosity, such that away from shocks the modelled flow is virtually inviscid.

Section 2 describes SPH artificial viscosity and previous efforts to reduce its adverse effects, while our new method is outlined in Section 3. The ability of the new scheme to reduce artificial viscosity but also to capture shocks is demonstrated in Sections 4 and 5, respectively. Finally, Section 6 concludes our study.

2 REDUCING UNWANTED ARTIFICIAL VISCOSITY

2.1 Standard SPH artificial viscosity

The traditional form of artificial viscosity (e.g. Monaghan 1992) adds the following terms to the momentum and energy equations, allowing the conversion of kinetic energy into heat:

$$(\dot{v}_i)_{AV} = - \sum_j m_j \Pi_{ij} \nabla_i \bar{W}_{ij}, \quad (2a)$$

$$(\dot{u}_i)_{AV} = \frac{1}{2} \sum_j m_j \Pi_{ij} \mathbf{v}_{ij} \cdot \nabla_i \bar{W}_{ij}, \quad (2b)$$

with the average kernel

$$\bar{W}_{ij} = \frac{1}{2} (W(|\mathbf{x}_{ij}|, h_i) + W(|\mathbf{x}_{ij}|, h_j)). \quad (3)$$

Here, $\mathbf{x}_{ij} \equiv \mathbf{x}_i - \mathbf{x}_j$ and $\mathbf{v}_{ij} \equiv \mathbf{v}_i - \mathbf{v}_j$, while h_i is the individual adaptive smoothing length of each SPH particle (for details on how h_i is adapted see Appendix A1). The artificial viscosity term is given by (Gingold & Monaghan 1983)

$$\Pi_{ij} = \begin{cases} \frac{-\alpha \bar{c}_{ij} \mu_{ij} + \beta \mu_{ij}^2}{\bar{\rho}_{ij}} & \text{for } \mathbf{v}_{ij} \cdot \mathbf{x}_{ij} < 0, \\ 0 & \text{otherwise,} \end{cases} \quad (4)$$

with

$$\mu_{ij} = \frac{\bar{h}_{ij} \mathbf{v}_{ij} \cdot \mathbf{x}_{ij}}{\mathbf{x}_{ij}^2 + \epsilon^2} \quad (5)$$

($\bar{h}_{ij} \equiv [h_i + h_j]/2$ and likewise for the average sound speed \bar{c}_{ij} and estimated density $\bar{\rho}_{ij}$). Since $\Pi_{ij} = 0$ for receding particle pairs, artificial viscosity does not affect expanding flows. This functional form of SPH artificial viscosity may seem rather ad hoc, but it is reasonably well motivated and emerged as the most useful one amongst several methods (Gingold & Monaghan 1983). Moreover, it is equivalent to the form of dissipation implicit in Riemann solvers (Monaghan 1997).

By expanding density and velocity in a Taylor series around \mathbf{x}_i , it is straightforward to show that these terms correspond to both a shear and a bulk viscosity. More quantitatively, if one assumes that, other than in equation (4), artificial viscosity acts between approaching and receding neighbours and that $\beta = 0$, the corresponding shear and bulk viscosity coefficients are (e.g. Meglicki, Wickramasinghe & Bicknell 1993) $\eta = (1/2)\alpha\kappa h c \rho$ and $\zeta = (5/3)\eta$, respectively, where the factor κ is of order unity and depends on the functional form of the kernel. This implies that artificial viscosity decreases with increasing resolution (smaller h). Thus, a straightforward though expensive way to reduce unwanted dissipation is to increase the resolution. In fact, one motivation for reducing artificial viscosity is to avoid this purely numerical necessity for high resolution.

Most SPH applications to date use the above treatment with $\alpha = 1$. The widely used code GADGET-2 (Springel 2005) employs a fixed α chosen at the start of the simulation (though Dolag et al. 2005 have implemented into GADGET-2 the improved method described in Section 2.3). Clearly, in complex situations, where strong and weak shocks are present as well as converging flows, any choice for α is unsatisfactory, leading to bad treatment of strong shocks, overdamping of converging flows, or both.

2.2 Balsara's method

The purpose of artificial viscosity is to allow for entropy generation across shocks and to stop particle interpenetration. To this end, only bulk viscosity is required, but the inherent shear viscosity is unnecessary. What is worse, this shear viscosity may seriously compromise simulations of shear flows, such as in a differentially rotating gas disc. In an effort to reduce the resulting artificial angular momentum dissipation, Balsara (1995) proposed to multiply Π_{ij} with a reduction factor $f_{ij} = (f_i + f_j)/2$ with

$$f_i = \frac{|\nabla \cdot \mathbf{v}_i|}{|\nabla \cdot \mathbf{v}_i| + |\nabla \times \mathbf{v}_i|} \quad (6)$$

(with velocity divergence and curl estimated using the SPH kernel estimator). This term diminishes the effect of artificial viscosity whenever the vorticity dominates the convergence. However, this method only reduces (but does not eliminate) unwanted dissipation in the presence of a rotating shear flow.

2.3 The method of Morris & Monaghan

Standard SPH artificial viscosity acts whenever the flow of the fluid converges, even if only weakly. For example, when a pulsating star contracts artificial viscosity damps its pulsation. Exactly the same happens to ordinary sound waves: standard SPH viscosity damps them, as demonstrated in Fig. 2, the faster the shorter the wavelength (because these are more poorly resolved).




Figure 2. As Fig. 1, but for SPH with standard ($\alpha = 1$) or Morris & Monaghan (1997) artificial viscosity, as well as our new method (only every fifth particle is plotted). Also shown are the undamped wave (solid) and lower amplitude sinusoids (dashed). Only with our method the wave propagates undamped, very much like SPH without any viscosity, as in Fig. 1.

With this in mind, Morris & Monaghan (1997) proposed to adapt the strength of artificial viscosity to the local convergence of the flow. To this end, they introduced the concept of individual adaptive viscosities α_i for each particle, replaced α in equation (4) by $\bar{\alpha}_{ij} = (\alpha_i + \alpha_j)/2$, and set $\beta \propto \bar{\alpha}_{ij}$. The individual viscosities are adapted according to the differential equation

$$\dot{\alpha}_i = (\alpha_{\min} - \alpha_i)/\tau_i + S_i, \quad (7)$$

with the velocity-based source term

$$S_i = \max\{-\nabla \cdot \mathbf{v}_i, 0\}, \quad (8)$$

and the decay time³

$$\tau_i = h_i/(2\ell c_i). \quad (9)$$

Here, $\alpha_{\min} = 0.1$ constitutes a lower limit for the artificial viscosity such that $\alpha_i = \alpha_{\min}$ for non-convergent flows. For a convergent flow, on the other hand, α_i grows above that value, guaranteeing the proper treatment of shocks. In the post-shock region, the flow is no longer convergent and α_i decays back to α_{\min} on the time-scale τ_i (typically $\ell = 0.1$ – 0.2). This method reduces the artificial viscosity away from shocks by an order of magnitude compared to standard SPH and gives equally accurate post- and pre-shock solutions (Morris & Monaghan 1997).

More recently, Rosswog et al. (2000) proposed to alter the adaptation equation (7) to⁴

$$\dot{\alpha}_i = (\alpha_{\min} - \alpha_i)/\tau_i + (\alpha_{\max} - \alpha_i) S_i, \quad (10)$$

with $\alpha_{\max} = 1.5$, while Price (2004) advocated $\alpha_{\max} = 2$. The effect of this alteration is first to prevent α_i to exceed α_{\max} and second to increase $\dot{\alpha}_i$ for small α_i , which ensures a faster viscosity growth, resulting in somewhat better treatment of shocks (Price 2004). This method may also be combined with the Balsara switch by applying the reduction factor (6) either to Π_{ij} (Rosswog et al. 2000) or to S_i (Morris & Monaghan 1997; Wetzstein et al. 2009).

The scheme of equations (8), (9) and (10) with $\alpha_{\min} = 0.1$, $\alpha_{\max} = 2$ and $\ell = 0.1$ is the current state of the art for SPH and is implemented in the codes PHANTOM (by Daniel Price) and VINE (Wetzstein et al. 2009). In Sections 4 and 5, we will frequently compare our novel scheme (to be described below) with this method and refer to

³ The factor 2 in the denominator of equation (9) accounts for the difference in the definition of the smoothing length h between us and Morris & Monaghan (1997).

⁴ This is equivalent to keeping (7) but multiplying the source term (8) by $(\alpha_{\max} - \alpha)$, which is what Rosswog et al. actually did.

it as the ‘M&M method’ or the ‘Price (2004) version of the M&M method’ as opposed to the ‘original M&M method’, which uses equation (7) instead of equation (10).

2.4 Critique of the M&M method

The M&M method certainly constitutes a large improvement over standard SPH, but low-viscosity flows, typical for many astrophysical fluids, are still inadequately modelled. After studying this and related methods in detail, we identify the following problems.

First, any $\alpha_{\min} > 0$ results in unwanted dissipation, for example of sound waves (see Fig. 2) or stellar pulsations (see Section 4.4), yet the M&M method requires $\alpha_{\min} \approx 0.1$. This necessity has been established by numerous tests (most notably of Price 2004) and is understood to originate from the requirement to ‘maintain order amongst the particles away from shocks’ (Morris & Monaghan 1997).

Secondly, there is a delay between the peak in the viscosity α and the shock front (see Fig. 3): the particle viscosities are still rising when the shock arrives. One reason for this lag is that integrating the differential equation (10) increases α_i too slowly: the asymptotic value

$$\alpha_s = \frac{\alpha_{\min} + \alpha_{\max} S_i \tau_i}{1 + S_i \tau_i} \quad (11)$$

is hardly ever reached before the shock arrives (and S_i decreases).

Thirdly, the source term (8) does not distinguish between pre- and post-shock regions: for a symmetrically smoothed shock it peaks at the exact shock position (in practice the peak occurs one particle separation in front of the shock; Morris & Monaghan 1997, see also Fig. 3). However, immediately behind the shock (or more precisely the minimum of $\nabla \cdot \mathbf{v}$), the (smoothed) flow is still converging and hence α continues to increase without need. A further problem is the inability of the source term (8) to distinguish between velocity discontinuities and convergent flows.

Finally, in strong shear flows the estimation of the velocity divergence $\nabla \cdot \mathbf{v}$, needed in (8), often suffers from substantial errors (see Appendix B1 for the reason), driving artificial viscosity without need. This especially compromises simulations of differentially rotating discs even when using the Balsara switch.

3 A NOVEL ARTIFICIAL VISCOSITY SCHEME

Our aim is a method which overcomes all the issues identified in Section 2.4 and in particular gives $\alpha_i \rightarrow 0$ away from shocks. To this end, we introduce a new shock indicator in Section 3.1, a novel technique for adapting α_i in Section 3.2 and a method to suppress false compression detections due to the presence of strong shear in Section 3.3.

3.1 A novel shock indicator

We need a shock indicator which not only distinguishes shocks from convergent flows, but, unlike $\nabla \cdot \mathbf{v}$, also discriminates between pre- and post-shock regions. This requires (at least) a second-order derivative of the flow velocity and we found the total time derivative of the velocity divergence, $\dot{\nabla} \cdot \mathbf{v} \equiv d(\nabla \cdot \mathbf{v})/dt$, to be most useful. As is evident from differentiating the continuity equation,

$$-\dot{\nabla} \cdot \mathbf{v} = d^2 \ln \rho / dt^2, \quad (12)$$

$\dot{\nabla} \cdot \mathbf{v} < 0$ indicates a non-linear density increase and a steepening of the flow convergence, as is typical for any pre-shock region.

Conversely, in the post-shock region $\nabla \cdot \mathbf{v} > 0$. This suggests to consider only negative values and, in analogy with equation (8), we define the new shock indicator

$$A_i = \xi_i \max\{-\nabla \cdot \mathbf{v}_i, 0\}. \quad (13)$$

Here, ξ_i is a limiter, detailed in Section 3.3, aimed at suppressing false detections of compressive flows in multidimensional flows.

3.2 Adapting individual viscosities

Instead of increasing α_i by integrating a differential equation, we set α_i directly to an appropriate local value $\alpha_{\text{loc},i}$ whenever this exceeds the current value for α_i . After extensive experimenting, we settled on the following simple functional form

$$\alpha_{\text{loc},i} = \alpha_{\text{max}} \frac{h_i^2 A_i}{v_{\text{sig},i}^2 + h_i^2 A_i}, \quad (14)$$

with the signal velocity⁵

$$v_{\text{sig},i} = \max_{|\mathbf{x}_{ij}| \leq h_i} \{\bar{c}_{ij} - \min\{0, \mathbf{v}_{ij} \cdot \hat{\mathbf{x}}_{ij}\}\}. \quad (15)$$

At the moment of passing through a shock (more precisely through a maximum of the flow convergence), A and hence α_{loc} return to zero and whenever $\alpha_i > \alpha_{\text{loc},i}$ we let α_i decay according to

$$\dot{\alpha}_i = (\alpha_{\text{loc},i} - \alpha_i)/\tau_i, \quad \tau_i = h_i/2\ell v_{\text{sig},i}. \quad (16)$$

We use $v_{\text{sig},i}$ rather than c in the decay time τ_i for internal consistency (this is of little practical relevance as $v_{\text{sig}} \approx c$ in the post-shock region). We use $\ell = 0.05$, such that the viscosity decays twice as slowly as in previous methods, avoiding some occasional minor post-shock ringing not present in methods with $\alpha_{\text{min}} > 0$. However, the traditional $\ell = 0.1$ also gives satisfactory results for most of our test problems.

3.3 Avoiding false compression detections

As explained in detail in Appendix B1, in multidimensional flows strong shear induces false detections of $\nabla \cdot \mathbf{v}$ with the standard SPH estimator even in the absence of particle disorder (noise). As shown in Appendix B2, these errors can be reduced by first estimating the velocity gradient matrix $\mathbf{V} \equiv \nabla \mathbf{v}$ and then obtaining $\nabla \cdot \mathbf{v}$ as its trace (we employ a similar method to estimate $\nabla \cdot \mathbf{v}$, see Appendix B3).

Unfortunately, even with this improved method false detections for $\nabla \cdot \mathbf{v}$ and $\nabla \cdot \mathbf{v}$ remain, for example in the situation of a differentially rotating disc. These still induce artificial viscosity, which may be significant in particular if c_s/h is small compared to the shear. The limiter ξ_i in equation (13) is aimed at suppressing such false detections by $\xi_i \rightarrow 0$ whenever the shear is much stronger than the convergence and no shock is present.

Having obtained the velocity gradient matrix \mathbf{V} , the shear is easily obtained as its traceless symmetric part $\mathbf{S} \equiv (\mathbf{V} + \mathbf{V}^T)/2 - v^{-1}(\nabla \cdot \mathbf{v})\mathbf{I}$ (with v the number of spatial dimensions), while the presence of a shock is indicated by

$$-1 \approx R_i \equiv \frac{1}{\hat{\rho}_i} \sum_j \text{sign}(\nabla \cdot \mathbf{v}_j) m_j W(|\mathbf{x}_i - \mathbf{x}_j|, h_i), \quad (17)$$

⁵ Various definitions for the signal velocity can be found in the SPH literature. Ours reflects the maximum velocity with which information can be transported between particles, but avoids $v_{\text{sig},i} \leq 0$.

since near a shock $\nabla \cdot \mathbf{v} < 0$ for all particles. After some experimenting, we found the following functional form for the limiter suitable:

$$\xi_i = \frac{|2(1 - R_i)^4 \nabla \cdot \mathbf{v}_i|^2}{|2(1 - R_i)^4 \nabla \cdot \mathbf{v}_i|^2 + \text{tr}(\mathbf{S}_i \cdot \mathbf{S}_i)}. \quad (18)$$

This functional form is similar to the Balsara limiter (6) in that it compares the flow convergence to a measure of the traceless part of the velocity gradient (the shear or the vorticity).

Alternatively, if one can be sure that no strong shear flows occur during the simulation, one may use the standard SPH estimator for $\nabla \cdot \mathbf{v}$ and estimate $\nabla \cdot \mathbf{v}$ from its change over the last time-step. However, the limiter is still desirable and one may use $|\nabla \times \mathbf{v}|^2$ instead of $\text{tr}(\mathbf{S} \cdot \mathbf{S}^T)$ in equation (18). We do not use this simplified version in the tests presented below, but our experiments indicate that such a method would pass all our tests except those of Sections 4.3 and 5.3, both involving strong shear.

3.4 Behaviour in typical situations

Before considering 2D and 3D test problems, we now assess the behaviour of our novel scheme, as well as that of the M&M method, in simple yet typical situations.

First, consider a sound wave of velocity amplitude $v_s \ll c$ and wave number $k \ll h^{-1}$ as example of a well-resolved weakly convergent flow. In this case, $A \simeq k^2 c v_s$ and $S \simeq k v_s$ at their respective maxima. Since $v_{\text{sig}} \simeq c \gg v_s$ we have $\alpha_{\text{loc}} \simeq \alpha_{\text{max}} h^2 k^2 (v_s/c)$, while for the M&M method the asymptotic value $\alpha_s \simeq \alpha_{\text{min}} + \alpha_{\text{max}} h k (v_s/c)/2\ell$. In the limit $kh \rightarrow 0$ of a well-resolved wave, $\alpha_{\text{loc}} \rightarrow 0$ faster than $\alpha_s \rightarrow \alpha_{\text{min}}$, such that even with $\alpha_{\text{min}} = 0$ the M&M method would be more viscous than our new scheme. Fig. 2 shows 1D sound-wave SPH runs, demonstrating that our new scheme behaves quasi-inviscid in this situation.

Following Morris & Monaghan (1997), we may also consider a simple homologous flow $\mathbf{v} = -a\mathbf{x}$ with $a > 0$, which approximates certain astrophysical problems involving collapse and does not require artificial dissipation. For this situation $S = 3a$ but $A = 0$ (a direct consequence of the ability of $\nabla \cdot \mathbf{v}$ to distinguish shocks from convergent flows), such that our new scheme remains inviscid, while the M&M method does not even for $\alpha_{\text{min}} = 0$.

Next, consider a strong shock with velocity discontinuity $\delta v \gg c$. Assuming that it is smoothed over one kernel width, we find maximum amplitudes of $S \simeq \delta v/h$ and $A \simeq (\delta v/h)^2$ (the exact values depend on the shock conditions and the functional form of the smoothing kernel). Since $v_{\text{sig}} \simeq h \nabla \cdot \mathbf{v} \sim \delta v$, our new scheme gives $\alpha_{\text{loc}} \sim \alpha_{\text{max}}$, while the asymptotic value (11) for the M&M method also approaches α_{max} .

While 3D simulations of strong shocks are presented in Section 5.2, Fig. 3 presents weak ram-shock simulations with $\delta v = 0.1c$ (top) and $\delta v = c$ (bottom) for our new scheme, the M&M method and standard SPH. In both regimes the peak in, respectively, α_{loc} and α_s is one particle farther in front of the shock with our new method than with the M&M method, which reflects the superiority of A over S to detect an incoming shock. This, combined with setting the viscosity parameter directly to the required value, results in the peak in α to occur two particle separations *before* the shock for our new method, while for the M&M method it peaks a similar length *behind* the shock.

With our new method, the viscous deceleration (bottom panels in Fig. 3) sets in about three particle separations before the weak and the strong shock, yielding good shock capturing properties in both cases. The M&M method, on the other hand, decelerates the flow

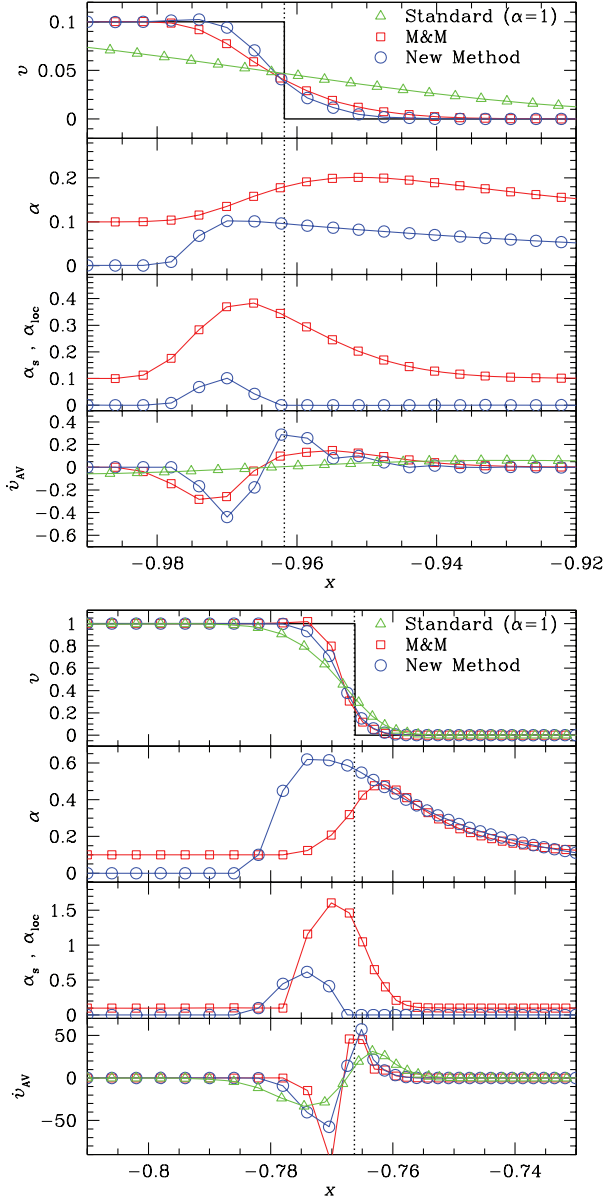


Figure 3. A 1D ram shock with $\delta v = 0.1$ (top) and $\delta v = 1$ (bottom) in ideal gas with $\gamma = 1.4$ simulated using standard SPH, the M&M and our new method. We compare the velocity, viscosity parameter, its asymptotic value and the viscous deceleration. Initially, the velocity is discontinuous with $v = -\delta v \text{ sign}(x)$, resulting in two shocks of δv propagating in either direction from the origin; the shock plotted propagates from right to left.

much earlier for a weak shock than for a strong shock and results in significant overdamping of weak shocks (which also pertains to density and internal energy – not shown in Fig. 3), while our method smooths both shocks over four particle separations (top panels in Fig. 3), the optimal SPH resolution in 1D. Note that standard SPH is hopeless: it oversmooths the strong shock and is completely incapable of dealing with the weak shock.

3.5 Maintaining particle order

The main point of our method is the absence of artificial viscosity away from shocks. Hence, if $\alpha_{\min} > 0$ was indeed required to maintain particle order, as previously argued in context of the M&M

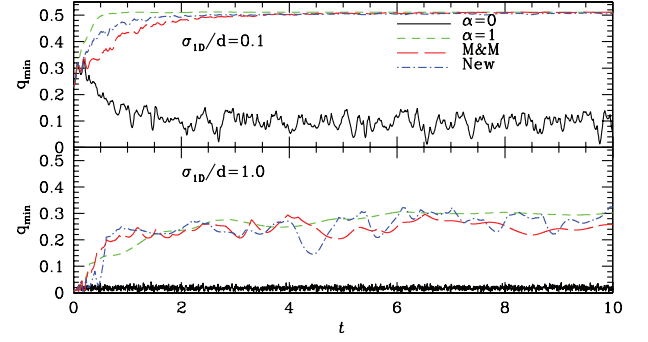


Figure 4. Time evolution of q_{\min} , defined in equation (19), for SPH simulations started from noisy initial conditions (see text). All SPH schemes with artificial viscosity suppress the noise equally well.

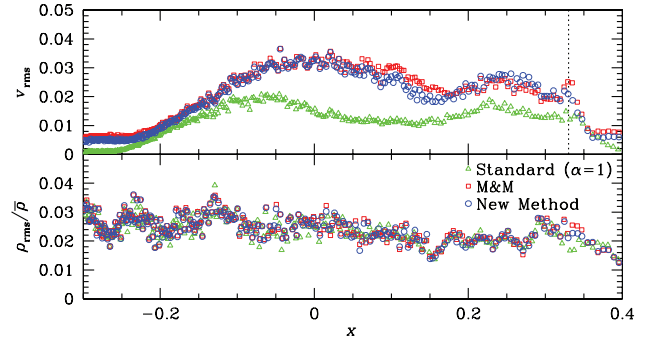


Figure 5. The rms amplitudes of density and velocity fluctuations for 3D simulations of the Sod (1978) shock tube test (see also Fig. 11). Initial conditions were prepared using a glass. The shock propagates to the right and is indicated by the dotted line; the velocity jump at the shock is 0.63.

method, our method should fail in this regard. Noise in SPH can emerge from shocks or carelessly generated initial conditions.

Let us first consider the time evolution of noisy initial conditions, generated by adding random displacements to particle positions representing noise-free hydrostatic equilibrium (the vertices of a face-centred-cubic grid, i.e. densest sphere packing). We consider two cases with the displacements in each dimension drawn from a normal distribution with rms amplitude equal to the nearest neighbour distance or a tenth of it, respectively. The time evolution of such noisy initial conditions can be distinguished by monitoring

$$q_{\min} \equiv \min_{i,j} \{|\mathbf{x}_{ij}|/h_i\}. \quad (19)$$

There are three possible scenarios. Either the particles settle back close to the original grid (q_{\min} approaches its grid value q_{grid}), form a glass (q_{\min} approaches a finite value $< q_{\text{grid}}$) or form dense clumps ('clumping instability', $q_{\min} \sim 0$). Fig. 4 plots the evolution of q_{\min} for $N_h = 40$ SPH neighbours (see also Appendix A1) when $q_{\text{grid}} \approx 0.529$. Clumping only occurs when $\alpha \equiv 0$, while for any viscous scheme tested the particles settle back on to the grid or form a glass with roughly similar time evolutions.

Post-shock noise occurs because the shock-induced compression disrupts the original particle order, but other than in the above test the viscosity is already switched on. In Fig. 5, we plot the amplitudes of the velocity and density noise in 3D simulations of the standard Sod (1978) shock tube test (see also Section 5.1). The three methods have similar levels of density noise, but standard SPH is less noisy in the velocities, which is not surprising given its stronger viscosity. However, between the two viscosity suppressing schemes there is little difference, even though $\alpha_{\min} = 0$ for our method. Similar




Figure 6. Steepening of a 1D sound wave: velocity and viscosity parameter versus position for standard SPH, the M&M method, our new scheme and GPH of first and second order (Cha & Whitworth 2003), each using 100 particles per wavelength. The solid curve in the top panel is the solution obtained with a high-resolution grid code.

results obtain for other shock tests and we conclude that our method is no worse than M&M's for maintaining particle order.

4 VISCOSITY SUPPRESSION TESTS

We now present some tests of low-Mach-number flows, where previous methods give too much unwanted dissipation.

4.1 Sound-wave steepening

The steepening of sound waves is a simple example demonstrating the importance of distinguishing between converging flows and shocks. As the wave propagates, adiabatic density and pressure oscillations result in variations of the sound speed, such that the density peak of the wave travels faster than the trough, eventually trying to overtake it and forming a shock.

In our test, a 1D sound wave with a velocity amplitude 10 per cent of the sound speed is used (ideal gas with $\gamma = 1.4$). Fig. 6 compares the velocity field at the moment of wave steepening for various SPH schemes, each using 100 particles, with a high-resolution grid simulation. The new method resolves the shock better than the M&M scheme, let alone standard SPH.

In Fig. 6, we also show results from Godunov-type particle hydrodynamics (GPH; Cha & Whitworth 2003), which differs from SPH by using the pressure P^* , found by solving the Riemann problem between particle neighbours, in the momentum and energy equations and avoids the need for explicit artificial viscosity. This substitution does not affect the energy or momentum conservation (Cha 2002), and indeed we find that both are well conserved. While the first-order GPH scheme is comparable to standard SPH and also to a Eulerian–Godunov grid code using the same Riemann solver without interpolation (not shown), the second-order GPH scheme resolves the discontinuity almost as well as our novel method.

4.2 1D converging flow test

Similar to sound-wave steepening, this test requires good treatment of convergent flows and weak shocks. The initial conditions are

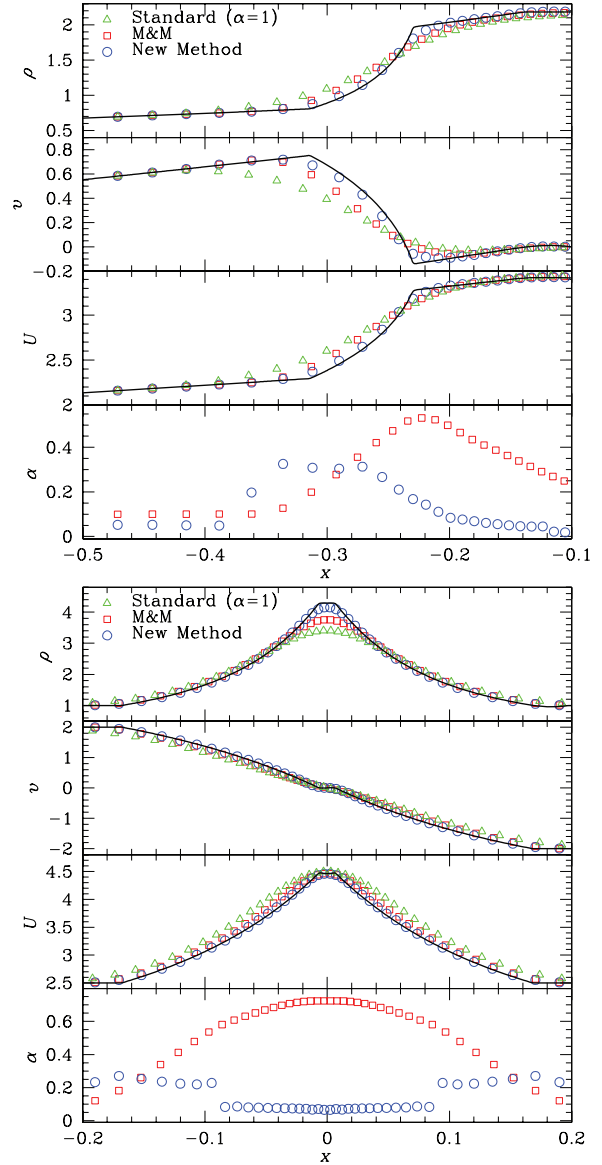


Figure 7. A 1D converging flow test with initially constant density and pressure and velocities given by equation (20) using an adiabatic equation of state with $\gamma = 1.4$. Top: run for $v_a = 1$ at $t = 0.3$; bottom: run for $v_a = 2$ at $t = 0.1$. The solid lines are the result of a high-resolution Eulerian grid-code simulation.

uniform pressure and density and a continuous flow velocity:

$$v = \begin{cases} 4(1+x)v_a & -1.00 < x < -0.75, \\ v_a & -0.75 < x < -0.25, \\ -4xv_a & -0.25 < x < 0.25, \\ -v_a & 0.25 < x < 0.75, \\ 4(1-x)v_a & 0.75 < x < 1.00. \end{cases} \quad (20)$$

As there is no analytical solution, we compare the results to a high-resolution grid-code simulation. We run tests for $v_a = 1$ and $v_a = 2$ as shown in the top and bottom panels of Fig. 7.

While the M&M switch certainly improves upon standard SPH, it still oversmooths the velocity profile as the viscosity is increased before a shock has formed. This is particularly evident in the velocity profile of the $v_a = 2$ case (bottom) near $x = 0$. The new switch keeps the viscosity low, in the $v_a = 2$ case an order of magnitude lower

than the M&M method. In fact, the agreement between our method and the high-resolution grid code is as good as one can possibly expect at the given resolution, in particular the velocity plateau and density amplitude around $x = 0$ in the $v_a = 2$ case (bottom) are correctly modelled.

4.3 2D Keplerian-ring test

In this test, a gaseous ring orbits a central point mass, neglecting the self-gravity of the gas. Initially, the ring is in equilibrium: pressure forces, attraction by the point mass and centrifugal forces balance each other. The Keplerian differential rotation implies that the flow is shearing and any viscosity causes the ring to spread (Lynden-Bell & Pringle 1974). This is indeed what Maddison, Murray & Monaghan (1996) found in SPH simulations without pressure forces.

Maddison et al. also found an instability to develop from the inner edge, which quickly breaks up the ring. They argue convincingly that this is the viscous instability (Lyubarskij, Postnov & Prokhorov 1994), which causes eccentric orbits at the inner edge of the ring to become more eccentric due to the viscous deceleration peaking at apocentre.

Imaeda & Inutsuka (2002) performed SPH simulations of the same problem but including pressure forces. They find a similar break-up of the ring after only few rotations and blame it on an inadequacy of the SPH scheme itself. We strongly suspect that Imaeda & Inutsuka encountered a form of the clumping instability, which appears to be particularly strong in 2D simulations of strong shear flows (though it may have been a dynamical instability inherent to gaseous Keplerian rings, e.g. Papaloizou & Pringle 1984, 1985; Goldreich & Narayan 1985). This numerical instability grows on a local hydrodynamical time and may therefore be suppressed by choosing the sound speed c much lower than the rotation speed v_ϕ . Indeed, Price (2004) and Monaghan (2006), who repeated these

and similar experiments with a very low sound speed, found no such numerical instabilities. A detailed investigation of these issues is clearly beyond the scope of our study and we merely compare our new scheme to previous methods for pressure forces with $c \ll v_\phi$ when the viscous instability should strike after few rotations depending on the strength of the artificial viscosity.

In our test, $GM = 1000$ for the central point mass, while the gas ring has Gaussian surface density centred on $r = 10$ with width (standard deviation) 2.5 represented by $N = 9745$ particles initially placed according to the method of Cartwright, Stamatellos & Whitworth (2009). This implies an orbital period of 2π and velocity of $v_\phi = 10$ at the ring centre. We choose a sound speed of $c = 0.01 \ll v_\phi$ to ensure that any dynamical instabilities of inviscid rings become important only after many periods.

Fig. 8 shows the particle distributions at various times for different SPH schemes. Only with our new method, the rings stay in their initial equilibrium configuration over at least five periods, while for the other methods, the inner parts of the ring soon become disordered leading to a catastrophic break-up after a few periods. It is noteworthy that this failure occurs despite the Balsara switch, which was designed specifically for applications like this.

Note that without the viscosity limiter ξ of equation (18), our novel method fails, precisely because of shear causing false detections of flow compression (as highlighted in Section 3.3 and Appendix B).

We also run similar tests with the central point mass replaced by a mass distribution (Plummer sphere or Kuzmin disc) with gravitational potential $\Phi = -GM/\sqrt{r^2 + s^2}$ with $s = 3$, such that the rotation curve of the disc also contains a rising part, similar to the situation in galactic discs. The outcome of these simulations (not shown) is essentially identical to that for the pure Keplerian rings: only our new method with viscosity limiter does not fall prey to the viscous instability.




Figure 8. Keplerian ring test: particle positions at various times for standard SPH with Balsara switch, the M&M method with and without Balsara switch and our new method without and with the viscosity limiter ξ of equation (18). Only for this last method the ring remains stable against a viscosity-induced instability. (Ring-like features at $r \lesssim 2$ are artefacts caused by the dynamical time close to the centre being short compared to the time-step.)

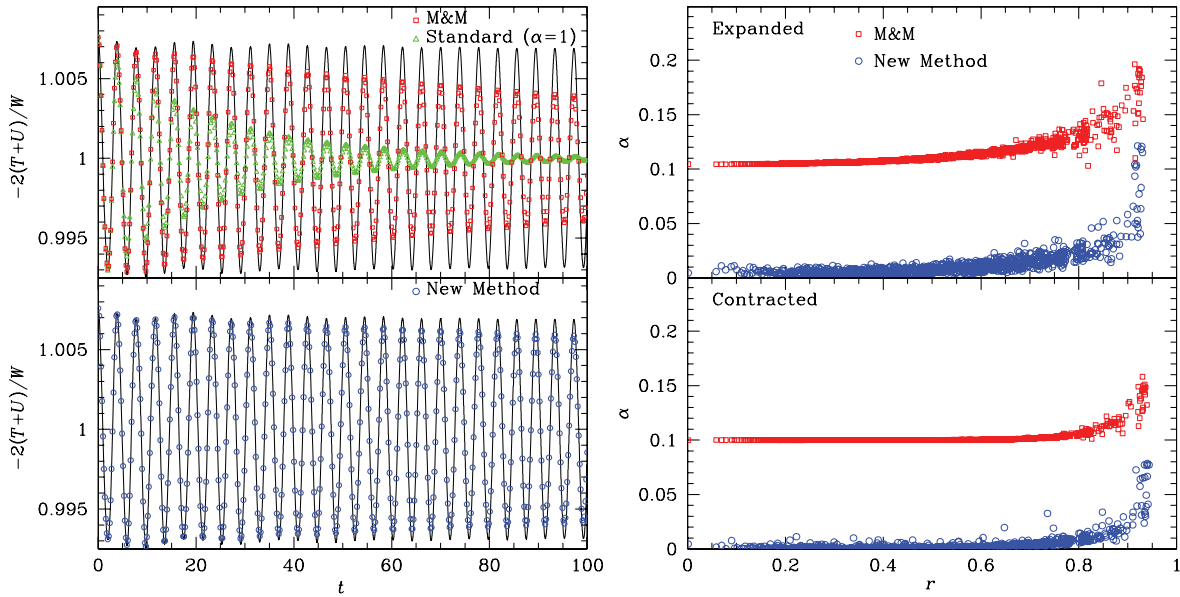


Figure 9. Left: virial ratio plotted versus time for SPH models of a radially oscillating polytrope which initially was in its fundamental radial eigenmode with relative radial amplitude of 0.01 and period 3.89. The solid curves are for a SPH model without any artificial viscosity. Right: the viscosity parameter α at $t = 97$ (maximum contraction) and $t = 99$ (maximum expansion) for every 100th particle. The new method keeps viscosity lower at the edge of the polytrope.

4.4 An oscillating polytropic sphere

The pulsations of a polytropic sphere are a good test for the adverse effects of artificial dissipation (Steinmetz & Müller 1993). We set up a polytropic sphere of 10^5 particles and induce oscillations in its fundamental mode (e.g. Cox 1980) with relative amplitude of 0.01 in radius and a period of $P = 3.8$.

In the absence of viscosity we expect the radial oscillations to continue with the initial amplitude and period over many oscillations. However, as with any numerical method some small amount of numerical dissipation may appear. None the less, such effects

should be small compared to the dissipation caused by artificial viscosity. Since the size of the radial perturbations increases with radius, we expect the oscillations to be small at the centre of the polytrope and therefore our new method to keep the viscosity low there. However, at the edge the size of the oscillations is more significant, and we may see an increase in viscosity at this point.

In order to track the oscillations, we monitor in Fig. 9 the time evolution of the virial ratio $-2(T + U)/W$ where T, U and W are the kinetic, internal and the gravitational energies, respectively. At maximum contraction the virial ratio is at its peak and at maximum expansion the virial ratio is lowest. With no artificial viscosity (solid

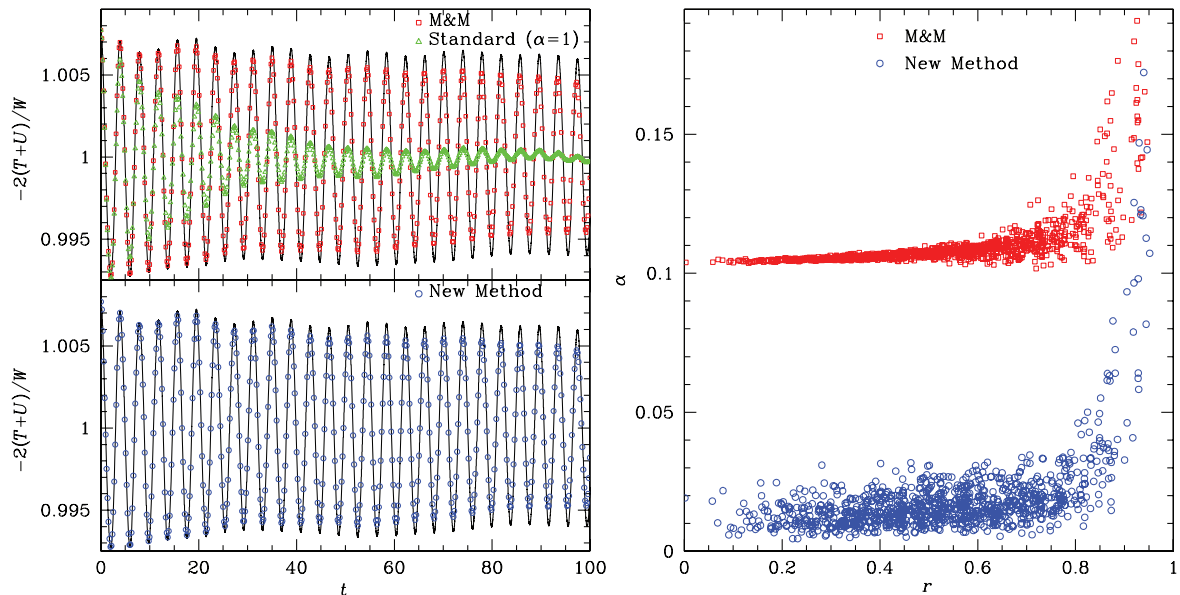


Figure 10. Same as Fig. 9, except that the sphere is in circular orbit around a point mass of 100 times its mass and with orbital radius 20 times its radius (the kinetic and potential energies are corrected for the contributions from the orbit). The viscosity parameter for every 100th particle is plotted at $t = 100$ (right).




Figure 11. Comparison of our new scheme and the M&M method for the standard Sod (1978) shock tube test with the analytic solution (solid).

curves in Fig. 9) the wave remains at constant amplitude barring a slight initial drop. The period averaged over 25 oscillations is $P = 3.89$, only slightly larger than the expected value. The reason for this discrepancy is most likely the unavoidable deviation of the (finite resolution) SPH model from a perfect polytropic sphere. This deviation also means that our SPH model is not exhibiting a pure eigenmode, but in addition contains some higher order modes at low amplitudes, resulting in some beating between them.

The M&M method results in a slow but continuous decay of the oscillations, though the period is hardly affected. This damping can be blamed largely on the finite α_{\min} (standard SPH damps the oscillation 10 times faster). Conversely, our new method, hardly damps the oscillations at all, because α is kept very small (except for the outermost layers where α is still below the M&M values).

We also run simulations where the oscillating polytropic sphere is on a circular orbit 20 times the radius of the sphere around a point mass 100 times that of the sphere (corresponding to a period of 56 time units). With this choice, the tidal radius is approximately four times the radius of the gas sphere, implying that tides are strong but not catastrophic. Since the orbital accelerations are much larger than those due to the polytropic oscillations, this is a tough test for any numerical scheme. In particular, Eulerian methods should have severe problems (this does exclude using corotating coordinates, which do not allow for tidal evolution of the orbit and are unavailable for eccentric orbits).

The time evolution of the virial ratio and the viscosity parameter α are shown in Fig. 10 for the same viscosity schemes as for the isolated case in Fig. 9. First note that the undamped simulations (solid curves) behave differently from the isolated case, exhibiting variations and a slight decay, both of which are most likely caused by the tidal field. As to be expected for any Lagrangian scheme, both SPH methods perform very similar to the isolated case, because neither $\nabla \cdot \mathbf{v}$ nor $\hat{\nabla} \cdot \mathbf{v}$ are affected by the orbital acceleration.

5 SHOCK CAPTURING TESTS

In this section, we subject our method to situations where artificial viscosity is required, mainly high Mach number shocks, and our aim is to demonstrate that it performs at least as well as the M&M method.

5.1 Sod shock tube test

The Sod (1978) shock tube test is a standard test for any shock capturing method and consists of an initial discontinuity in pressure and density leading to the production of a rarefaction wave, contact discontinuity and shock wave, which forms from the steepening of a subsonic wave. The whole system is subsonic with a maximum Mach number of $\mathcal{M} \approx 0.63$ in the pre-shock region. We perform the test in 3D at a resolution of 200 particle layers in the high-density region.

The density, energy, velocity and viscosity for standard SPH as well as the M&M and our method are shown in Fig. 11. As for the 1D ram test (see Fig. 3), our new method switches on viscosity already in the pre-shock region peaking about one smoothing length before the actual shock front (which travels to the right in Fig. 11), whereas the M&M switch turns on viscosity later, lagging our method by about four particle separations. As a consequence, the transition of the fluid values across the shock front is slightly smoother with our method than with the M&M method.

Note that the irregularities around the contact discontinuity at $x = 0.138$ common to all schemes tested are not related to artificial viscosity (the irregularities in α at that point could be removed by choosing non-zero initial α at the initial discontinuity); they can be alleviated by artificial conductivity (Price 2004; Price 2008).

5.2 Strong shocks and particle penetration

In Section 3.4 and Fig. 3 we already demonstrated that our new method is superior to the M&M scheme in resolving ‘subsonic shocks’ (velocity discontinuities smaller than the sound speed) and comparable in resolving shocks of Mach number ~ 1 . Here, we extend this comparison to high Mach numbers. Fig. 12 shows the result for the 1D ram test with $\mathcal{M} = 50$. The Price (2004) version of the M&M method, which uses equation (10) with $\alpha_{\max} = 2$, is implemented in some contemporary SPH codes, and has been used in our tests so far, fails this test: α remains too low and as a consequence the velocity discontinuity is not correctly smoothed and some post-shock ringing occurs. To give credit to Morris & Monaghan (1997), we also tested their original method and find it to work well (stars in Fig. 12). Our new method works about




Figure 12. Same as Fig. 3 but for $\mathcal{M} = 50$. We distinguish between the original M&M method (using equation 7) and the Price (2004) version (using equation 10 with $\alpha_{\max} = 2$), which has been denoted ‘M&M’ in all figures so far.

as well as the original M&M scheme, with α reaching the same level, though our scheme detects the coming shock much earlier: α is ahead of the original M&M method by about four particle separations.

Whilst the main role of artificial viscosity is to resolve shocks by transferring entropy, a secondary but vital role is to prevent particle penetration, which requires strong viscosity in high Mach number shocks. Bate (1995) performed many tests to determine the value of the parameters α and β needed to prevent particle penetration in ram shock tests of various Mach numbers. For particles arranged in face-centred-cubic or cubic grids, Bate found that appropriate values for the viscosity parameters can prevent particle penetration for shocks up to $\mathcal{M} = 8$. Most SPH practitioners opt for a value of $\beta = 2\alpha$ (Morris & Monaghan 1997).

To determine the correct value of β required for the new scheme, we perform high-resolution 3D runs of ram shocks with $\mathcal{M} = 20$ and various values for β/α . We smooth the initial velocity discontinuity, as suggested by Monaghan (1997), to provide the method with a situation realistic for SPH, such as would have arisen for a shock forming from continuous initial conditions.

For different values of β/α with our viscosity scheme and the two variants of the M&M switch, we plot in Fig. 13 the x and y positions (for all values of z) of particles near the shock front at a late time. The colour coding distinguishes particles which at that time should be upstream (red) or downstream (green). Our scheme prevents particle penetration with $\beta = \alpha$ (for $\beta = 0$ there is one layer of overlap). The original M&M scheme with the standard choice $\beta = 2\alpha$ also avoids particle penetration, but not the Price (2004) version, again a consequence of too little viscosity.

5.3 A shearing shock

This test combines a shock with a perpendicular shear and presents a difficult test for any SPH scheme. We use periodic boundary conditions and start from a face-centred-cubic grid and velocities:

$$v_x = -\delta v \operatorname{sign}(x), \quad v_y = s \sin(\pi x) \quad \text{and} \quad v_z = 0. \quad (21)$$

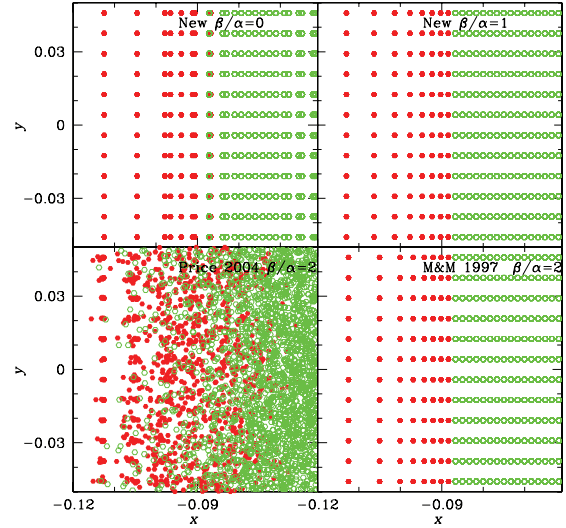


Figure 13. Particle positions in the x - y plane of 3D simulations of a $\mathcal{M} = 20$ ram shock along x direction. Particles are coloured red if their initial positions was $x_0 < -0.45$ and green if $x_0 > -0.45$.




Figure 14. Shearing shock test: density and velocity for various SPH schemes (symbols) and a grid-code simulation (curve). Initial velocities are given by equation (21) with $s = 100\delta v = 100c$.

In Fig. 14, we present results for various SPH simulations as well as a grid-code run for $s = 100\delta v = 100c$. The M&M method produces a large viscosity due to the shear-induced errors in $\nabla \cdot \mathbf{v}$, leading to spurious results. Using the Balsara limiter with either M&M or Standard SPH gives in much better results, though the shock is clearly oversmoothed. The new scheme is able to limit the viscosity to the correct level, allowing good capturing of the shock and retaining particle order in the post-shock region.

Note that this is a difficult test for any SPH implementation: without viscosity reduction (as in standard SPH) the shear flow is strongly damped, while viscosity reduction schemes (M&M as well as ours) suffer from the problem of shear-induced errors. These potentially result in too much viscosity and oversmoothing of the shock. Our limiter was able to control this problem, but for yet larger ratios $s/\delta v$ of shear to shock amplitude this problem becomes too difficult for any SPH implementation.

5.4 Evrard test

In this test the inward gravitational pull of a gas cloud exceeds its outward pressure force causing the cloud to collapse under its own self-gravity. The initial conditions consist of a gas sphere with density profile (Evrard 1988)

$$\rho(r) = \frac{M}{2\pi R^2} \frac{1}{r} \quad (22)$$

for $r < R$ and $\rho = 0$ for $r > R$. Initially the gas is at rest and has constant specific internal energy $u = 0.05 GM/R$, which corresponds to a virial ratio $-2U/W = 0.075 \ll 1$. The initial gravitational inward pull is the same at each radius, while the pressure forces decline outwards, leading to collapse and, as a consequence, formation of a shock, which steepens and evolves into a strong shock propagating outwards as more incoming material joins the jam. Even though the problem is initially spherically symmetric, the SPH realization of initial conditions cannot be exactly spherically symmetric and the system may well evolve away from sphericity, for instance driven by dynamical instabilities.

We use a unit system such that $G = R = M = 1$ and represent the cloud by 100 280 SPH particles, initially placed on a face-centred-cubic grid which is then radially stretched to match the density.

Fig. 15 compares the simulation results for our method, the original M&M method and a 1D calculation by Steinmetz & Müller (1993) using the piece-wise parabolic method (PPM).

At early times ($t = 0.39$, left-hand column) the results from all three methods match very well, but the M&M scheme already shows a large viscosity. At later times a shock forms (at $r \approx 0.13$ by $t = 0.78$), which moves outwards until it reaches the end of the sphere, when a significant fraction of the gas still has outwards velocities (by $t = 1.95$). The most obvious difference between the two SPH schemes is the amount of (artificial) dissipation: the M&M method is much more viscous, resulting in significant oversmoothing of the shock front by $t = 0.78$ accompanied by unphysical pre-shock heating as visible in the entropy (K) profile. Our new scheme agrees better with the 1D calculation, in particular in the inner (post-shock) regions. Note that with our new method α peaks well before the shock arrives (at $t = 1.17$), while for the M&M method the peak in α appears actually slightly after the shock.

We found this a valuable test as early versions of our scheme tended to be far too viscous, while our final version passes this test ahead of the M&M switch. Standard SPH (not shown in the figure) shows similar results, though the shock at $t = 0.78$ appears less smoothed than with the M&M method but more smoothed than with the new scheme.




Figure 15. The Evrard test (see text for the initial set-up): shown are various physical quantities ($K = P\rho^{-\gamma}$ is the entropy function) and α at different times for SPH simulations with $N = 10^5$ particles using our new viscosity scheme (blue) or the original M&M method (red). Also shown (black) are the results from 1D PPM calculation (Steinmetz & Müller 1993). Not every particle is plotted.

6 SUMMARY

Any hydrodynamical numerical method requires some form of artificial viscosity in order to resolve shocks (in grid methods, artificial viscosity is implicit in the Riemann solver; Monaghan 1997). In grid codes, such as RAMSES (Teyssier 2002), interpolation methods are employed to effectively suppress artificial viscosity away from shocks. Most SPH simulations to this date hardly use such precautions and, as a consequence, adiabatic oscillations and shear flows are damped. Note that this affects state-of-the-art simulations of e.g. galaxy formation, which usually only employ Balsara's (1995) rather inefficient method to reduce some adverse effects of artificial viscosity on rotation discs.

The method of Morris & Monaghan (1997), which reduces the default amount of artificial viscosity by an order of magnitude compared to standard SPH practice, has only recently been recognized as advantageous. In this method, explained in detail in Section 2.3, individual artificial viscosities α_i are adapted by integrating a differential equation. Though constituting a major improvement, this method remains unsatisfactory, because it still damps adiabatic oscillations and oversmooths weak shocks, as we argued in Section 3 and demonstrated in Section 4.

In Section 3, we present a novel method, which improves upon that of Morris & Monaghan in four important ways.

- (i) We set $\alpha_{\min} = 0$ enabling $\alpha_i \rightarrow 0$ away from shocks and effectively modelling the fluid as inviscid.
- (ii) We use $\dot{\nabla} \cdot \mathbf{v} \equiv d(\nabla \cdot \mathbf{v})/dt < 0$ rather than $\nabla \cdot \mathbf{v} < 0$ as shock indicator. This distinguishes pre-shock from post-shock regions (where $\dot{\nabla} \cdot \mathbf{v} > 0$ but $\nabla \cdot \mathbf{v} < 0$) and discriminates much better between converging flows and weak shocks.
- (iii) We set α_i directly to an appropriate local value α_{loc} , instead of growing it by integrating a differential equation.
- (iv) We use an improved estimator for $\nabla \cdot \mathbf{v}$ and $\dot{\nabla} \cdot \mathbf{v}$ and employ a limiter to avoid viscosity driven by shear-induced errors.

Together these novelties result in a significantly improved artificial viscosity method, in particular the viscosity is increased to an appropriate level well before an incoming shock. The implementation details, i.e. the precise way of setting α_{loc} from $\dot{\nabla} \cdot \mathbf{v}$ and the exact form of the limiter, may well be subject to improvements. Any reader who considers modifying these details is advised to consider the behaviour of the resulting method for a test suite comprising noise suppression as well as shear and strong shocks, for example the tests of Figs 4, 8 and 14.

For static equilibria $\nabla \cdot \mathbf{v} = 0$ and $\dot{\mathbf{v}} = 0$, and our new shock indicator (as well as the M&M shock indicator) are only triggered by velocity noise. As long as particle order is maintained, such noise triggers only negligible amounts of viscosity, unlike the situation with the M&M method, whose minimum viscosity $\alpha_{\min} = 0.1$ is often sufficient to affect the simulations (as demonstrated in Section 4). None the less, the noise-induced viscosity is sufficient to suppress particle disorder, as demonstrated in Section 3.5.

For dynamic equilibria $\nabla \cdot \mathbf{v} = 0$ (and $\dot{\nabla} \cdot \mathbf{v} = 0$) but $\dot{\mathbf{v}} \neq 0$. However, in multidimensional flows strong shear induces false detections of $\nabla \cdot \mathbf{v}$ (and $\dot{\nabla} \cdot \mathbf{v}$), even with best possible particle order, for reasons explained in Appendix B1. In simulations of differentially rotating discs, this problem strongly affects the M&M method (even when using the Balsara switch). We avoid this problem by applying a limiter (see Section 3.3) as well as using improved estimators for $\nabla \cdot \mathbf{v}$ and $\dot{\nabla} \cdot \mathbf{v}$, see Appendix B2 for details. (Alternatively, if no strong shear flows are present, the standard estimators should

suffice, though still in conjunction with a limiter using $|\nabla \times \mathbf{v}|$ as a proxy for the shear amplitude.)

These improved estimators also provide the full velocity and acceleration gradient matrices for each particle (and increase the computational costs by ~ 30 per cent). The knowledge of the velocity gradient matrix \mathbf{V} and its traceless symmetric part, the shear \mathbf{S} , is also useful for implementing physical viscosity

$$\rho \dot{\mathbf{v}} = \nabla \cdot [\eta \mathbf{S} + \zeta \text{tr}(\mathbf{V})] \quad (23)$$

(with η and ζ the shear and bulk viscosity coefficients) in SPH.

In Sections 3.5, 4 and 5, we demonstrate convincingly that our technique successfully deals with the following four situations:

shocks are resolved at least as well, if not better, than with any previous technique;

adiabatic oscillations, such as sound waves or stellar pulsations, remain undamped over many periods, which was not possible with any previous SPH implementation;

strong shear flows, such as in accretion discs, are modelled virtually inviscid, while shearing shocks are well resolved without being oversmoothed;

particle disorder is suppressed at least as well as with the M&M method.

In particular, in the regime of convergent flows and weak shocks our new method is far superior to any previous scheme, which all required a significant increase in resolution just to suppress adverse effects of artificial viscosity.

ACKNOWLEDGMENTS

Research in theoretical astrophysics at Leicester is supported by a STFC rolling grant. We thank all members of the Leicester Theoretical Astrophysics group, in particular Graham Wynn and Justin Read, for helpful discussions, Seung-Hoon Cha for contributing the GPH data to Fig. 6, Matthias Steinmetz for providing the PPM data of the Evrard test for Fig. 15 and the anonymous referee for helpful comments and enhanced scrutiny.

REFERENCES

- Balsara D., 1995, *J. Comput. Phys.*, 121, 357
- Bate M., 1995, PhD thesis, Univ. Cambridge, UK
- Cartwright A., Stamatellos D., Whitworth A. P., 2009, *MNRAS*, 395, 2373
- Cha S.-H., 2002, PhD thesis, Cardiff Univ., UK
- Cha S.-H., Whitworth A. P., 2003, *MNRAS*, 340, 73
- Cox J. P., 1980, *Theory of Stellar Pulsation*. Princeton Univ. Press, Princeton, NJ
- Dehnen W., 2002, *J. Comput. Phys.*, 179, 27
- Dolag K., Vazza F., Brunetti G., Tormen G., 2005, *MNRAS*, 364, 753
- Evrard A. E., 1988, *MNRAS*, 235, 911
- Gingold R. A., Monaghan J. J., 1977, *MNRAS*, 181, 375
- Gingold R. A., Monaghan J. J., 1983, *MNRAS*, 204, 715
- Goldreich P., Narayan R., 1985, *MNRAS*, 213, 7p
- Imaeda Y., Inutsuka S.-i., 2002, *ApJ*, 569, 501
- Lucy L. B., 1977, *AJ*, 82, 1013
- Lynden-Bell D., Pringle J. E., 1974, *MNRAS*, 168, 603
- Lyubarskij Y. E., Postnov K. A., Prokhorov M. E., 1994, *MNRAS*, 266, 583
- Maddison S. T., Murray J. R., Monaghan J. J., 1996, *Publ. Astron. Soc. Australia*, 13, 66
- Meglicki Z., Wickramasinghe D., Bicknell G. V., 1993, *MNRAS*, 264, 691
- Monaghan J. J., 1992, *ARA&A*, 30, 543
- Monaghan J. J., 1997, *J. Comput. Phys.*, 136, 298
- Monaghan J. J., 2006, *MNRAS*, 365, 199
- Monaghan J. J., Lattanzio J. C., 1985, *A&A*, 149, 135

- Morris J. P., Monaghan J. J., 1997, *J. Comput. Phys.*, 136, 41
 Papaloizou J. C. B., Pringle J. E., 1984, *MNRAS*, 208, 721
 Papaloizou J. C. B., Pringle J. E., 1985, *MNRAS*, 213, 799
 Price D. J., 2004, PhD thesis, Univ. Cambridge, UK
 Price D. J., 2008, *J. Comput. Phys.*, 227, 10040
 Rosswog S., Davies M. B., Thielemann F.-K., Piran T., 2000, *A&A*, 360, 171
 Sod G. A., 1978, *J. Comput. Phys.*, 27, 1
 Springel V., 2005, *MNRAS*, 364, 1105
 Springel V., Hernquist L., 2002, *MNRAS*, 333, 649
 Steinmetz M., Müller E., 1993, *A&A*, 268, 391
 Teyssier R., 2002, *A&A*, 385, 337
 Wetzstein M., Nelson A. F., Naab T., Burkert A., 2009, *ApJS*, 184, 298

APPENDIX A: DETAILS OF THE SPH SCHEME

For completeness, we give here a brief description of our SPH method, which is largely similar to previous methods, but may differ in some details.

A1 Density and adaptive smoothing lengths

Let ν denote the number of spatial dimensions, then we adapt the individual smoothing lengths h_i such that $h_i^\nu \hat{\rho}_i = M_h$ with $M_h \equiv mN_h/V_\nu$ a global constant, defined in terms of the number N_h of neighbours, the mass m of each SPH particle and the volume V_ν of the unit sphere. In this work, we use $N_h = 5, 13$ and 40 for $\nu = 1, 2$ and 3 dimensions, respectively. Inserting the density estimator (1), we find

$$h_i^\nu \hat{\rho}_i = \sum_j m_j w(r_{ij}) \quad \text{with} \quad r_{ij} \equiv |\mathbf{x}_{ij}|/h_i, \quad (\text{A1})$$

where we have rewritten the SPH kernel as $W(|\mathbf{x}_{ij}|, h_i) = h_i^{-\nu} w(r_{ij})$ with the dimensionless function $w(r)$. For this work, we employ the usual cubic spline kernel (Monaghan & Lattanzio 1985):

$$w(r) = \binom{\nu+3}{3} \frac{1}{V_\nu(2-2^{-\nu})} \begin{cases} 1-6r^2(1-r) & r \leq 1/2, \\ 2(1-r)^3 & 1/2 < r < 1, \\ 0 & \text{otherwise.} \end{cases} \quad (\text{A2})$$

At each time-step, the h_i are adjusted by performing one Newton–Raphson step in $\log h - \log(h^\nu \hat{\rho})$ space, i.e.

$$h_i \leftarrow h_i \left(\frac{M_h}{h_i^\nu \hat{\rho}_i} \right)^{f_i/\nu}, \quad (\text{A3})$$

with a factor of order unity

$$f_i = -\nu \frac{\sum_j m_j w_{ij}}{\sum_j m_j r_{ij}^2 \tilde{w}_{ij}}, \quad (\text{A4})$$

where $w_{ij} \equiv w(r_{ij})$ and $\tilde{w}(r) \equiv w'(r)/r$. This method converges extremely well, except when h_i was too much small. In this case, faster convergence can be achieved by subtracting the self-contribution (which does not depend on h_i). Thus, whenever $h_i^\nu \hat{\rho}_i < M_h$ we use instead of (A3),

$$h_i \leftarrow h_i \left(\frac{M_h - m_i w(0)}{h_i^\nu \hat{\rho}_i - m_i w(0)} \right)^{\tilde{f}_i/\nu}, \quad \text{with} \quad \tilde{f}_i = -\nu \frac{\sum_{j \neq i} m_j w_{ij}}{\sum_{j \neq i} m_j r_{ij}^2 \tilde{w}_{ij}}. \quad (\text{A5})$$

The time derivatives \dot{h}_i are obtained by demanding $d(h_i^\nu \hat{\rho}_i)/dt = 0$:

$$\frac{\dot{h}_i}{h_i} = \frac{\sum_j m_j \mathbf{v}_{ij} \cdot \mathbf{x}_{ij} \tilde{w}_{ij}}{\sum_j m_j \mathbf{x}_{ij}^2 \tilde{w}_{ij}}. \quad (\text{A6})$$

A2 Pressure forces

We use SPH equations of motion derived from the simple SPH Lagrangian $\mathcal{L} = \sum_k m_k [\mathbf{x}_k^2/2 - u_k]$. Together with the relation⁶ $du/d\rho = P/\rho^2$, this gives

$$\dot{\mathbf{v}}_i = -\frac{1}{m_i} \frac{\partial \mathcal{L}}{\partial \mathbf{x}_i} = -\sum_j m_j \mathbf{x}_{ij} \left(\frac{P_i f_i}{\hat{\rho}_i^2 h_i^{\nu+2}} \tilde{w}_{ij} + \frac{P_j f_j}{\hat{\rho}_j^2 h_j^{\nu+2}} \tilde{w}_{ji} \right), \quad (\text{A7})$$

where the factors f_i and f_j (equation A4) arise from the fact that the derivatives $\partial \hat{\rho}_k / \partial \mathbf{x}_i$ have to be taken at fixed $h_k^\nu \hat{\rho}_k$. The work done by these pressure forces has to be balanced by

$$\dot{u}_i = -\nu \frac{P_i \dot{h}_i}{\hat{\rho}_i h_i} = \frac{P_i f_i}{\hat{\rho}_i^2 h_i^{\nu+2}} \sum_j m_j \mathbf{v}_{ij} \cdot \mathbf{x}_{ij} \tilde{w}_{ij}. \quad (\text{A8})$$

A3 Artificial viscosity

For the artificial viscosity drag and heating we actually use

$$(\dot{\mathbf{v}}_i)_{\text{AV}} = -\sum_j m_j \mathbf{x}_{ij} \frac{\tilde{\Pi}_{ij}}{2} \left(\frac{\alpha_i f_i}{\hat{\rho}_i h_i^{\nu+2}} \tilde{w}_{ij} + \frac{\alpha_j f_j}{\hat{\rho}_j h_j^{\nu+2}} \tilde{w}_{ji} \right), \quad (\text{A9})$$

$$(\dot{u}_i)_{\text{AV}} = \sum_j m_j \mathbf{v}_{ij} \cdot \mathbf{x}_{ij} \frac{\tilde{\Pi}_{ij}}{2} \frac{\alpha_i f_i}{\hat{\rho}_i h_i^{\nu+2}} \tilde{w}_{ij}, \quad (\text{A10})$$

with $\tilde{\Pi}_{ij} = -\mu_{ij}(\tilde{c}_{ij} - b\mu_{ij})$, where

$$\mu_{ij} = \begin{cases} \frac{2\mathbf{v}_{ij} \cdot \mathbf{x}_{ij}}{(r_{ij}^2 + r_{ji}^2) \tilde{h}_{ij}} & \text{for } \mathbf{v}_{ij} \cdot \mathbf{x}_{ij} < 0, \\ 0 & \text{otherwise,} \end{cases} \quad (\text{A11})$$

while the parameter b has the meaning of β/α for traditional SPH. Note that equations (A9) and (2a) differ only by $\mathcal{O}(\tilde{h}_{ij}^2)$. The difference between equations (A10) and (2b) is more pronounced since, similarly to equation (A8), we do not symmetrize the contributions with respect to i and j .

A4 Time integration

Our scheme employs a kick-drift-kick leapfrog time integrator, which is second-order accurate. With this scheme, a full (global) time-step of size δt consists of the following substeps (‘ \leftarrow ’ means ‘is replaced by’).

Initial kick. Compute \mathbf{v}_i and u_i at half step:

$$\begin{aligned} \tilde{\mathbf{v}}_i &= \mathbf{v}_i + \frac{1}{2} \delta t \dot{\mathbf{v}}_i, \\ \tilde{u}_i &= u_i + \frac{1}{2} \delta t \dot{u}_i. \end{aligned} \quad (\text{A12})$$

Full drift. Advance t and \mathbf{x}_i by a full step:

$$\begin{aligned} t &\leftarrow t + \delta t, \\ \mathbf{x}_i &\leftarrow \mathbf{x}_i + \delta t \tilde{\mathbf{v}}_i. \end{aligned} \quad (\text{A13})$$

⁶ Alternatively, for an ideal-gas equation of state one may replace u in the Lagrangian with $u = K \hat{\rho}^{\gamma-1}/(\gamma-1)$ and consider the entropy function $K = P \hat{\rho}^{-\gamma}$ to be constant (Springel & Hernquist 2002).

Prediction. Predict \mathbf{v}_i, u_i and h_i at full step:

$$\begin{aligned}\mathbf{v}_i &\leftarrow \mathbf{v}_i + \delta t \dot{\mathbf{v}}_i, \\ u_i &\leftarrow u_i \exp(\delta t \dot{u}_i / u_i), \\ h_i &\leftarrow h_i \exp(\delta t \dot{h}_i / h_i).\end{aligned}\quad (\text{A14})$$

Sweep 0. Compute $h_i^v \hat{\rho}_i$ and f_i (equations A1 and A4).

Adapt. Adjust h_i (equation A3 or A5).

Sweep 1. Compute $\hat{\rho}_i$, f_i and \dot{h}_i (equations A1, A4 and A6) as well as $\nabla \cdot \mathbf{v}_i$, $\dot{\mathbf{v}} \cdot \mathbf{v}_i$ and R_i (equations B8, B12 and 17, using $\dot{\mathbf{v}}$ and $\nabla \cdot \mathbf{v}$ from the previous time-step).

Between sweeps. Obtain P_i and c_i from $\hat{\rho}_i$ and u_i via the equation of state, and adapt α_i via (using equations 9 and 14)

$$\alpha_i \leftarrow \begin{cases} \alpha_{\text{loc}} & \text{if } \alpha_i < \alpha_{\text{loc}}, \\ \alpha_{\text{loc}} + (\alpha_i - \alpha_{\text{loc}}) \exp(-\delta t / \tau_i) & \text{otherwise.} \end{cases} \quad (\text{A15})$$

Sweep 2. Compute $\dot{\mathbf{v}}_i$ (equations A7 and A9 plus gravitational forces) and \dot{u}_i (equations A8 and A10 plus external heating or cooling).

Final kick. Set \mathbf{v}_i and u_i at full step:

$$\begin{aligned}\mathbf{v}_i &= \tilde{\mathbf{v}}_i + \frac{1}{2} \delta t \dot{\mathbf{v}}_i, \\ u_i &= \tilde{u}_i + \frac{1}{2} \delta t \dot{u}_i.\end{aligned}\quad (\text{A16})$$

In the initial kick and prediction steps, the time derivatives are known from the previous time-step (in case of the very first time-step, they need to be pre-computed). Note that the quantities predicted in (A14) enter the final \mathbf{v}_i and u_i only indirectly via the computation of the time derivatives.

We use an oct-tree, generated just before sweep 0, to find all interacting particle pairs, which are then remembered in an interaction list, whereby allowing for the fact that h_i may grow slightly during adjustment (just after sweep 0). Utilizing this interaction list in sweeps 1 and 2 is much faster than further tree walks. The same oct-tree is also used in computing gravitational forces, as outlined by Dehnen (2002).

Our scheme can also be implemented with adaptive individual time-steps organized in a hierarchical block-step scheme, though we have not used this in the tests presented in this study.

APPENDIX B: ESTIMATING $\nabla \cdot \mathbf{v}$ AND $\dot{\mathbf{v}} \cdot \mathbf{v}$

B1 Failure of the standard SPH estimator for $\nabla \cdot \mathbf{v}$

Our constraint that $h_i^v \hat{\rho}_i$ be constant (see Section A1) implies $\dot{\rho}_i / \hat{\rho}_i = -v \dot{h}_i / h_i$. Together with the continuity equation $\dot{\rho} + \rho \nabla \cdot \mathbf{v} = 0$ and equation (A6) this yields the simple velocity-divergence estimate:

$$\widehat{\nabla \cdot \mathbf{v}}_i = v \frac{\sum_j m_j \mathbf{v}_{ij} \cdot \mathbf{x}_{ij} \tilde{w}_{ij}}{\sum_j m_j \mathbf{x}_{ij}^2 \tilde{w}_{ij}}. \quad (\text{B1})$$

While this estimate satisfies the continuity equation for the SPH density estimate $\hat{\rho}_i$, it is not necessarily accurate. To see this, consider the matrix (\otimes denotes the outer or dyadic vector product)

$$\mathbf{D}_i = \sum_j \mathbf{v}_{ij} \otimes \mathbf{x}_{ij} \tilde{w}_{ij}, \quad (\text{B2})$$

with \tilde{w}_{ij} some weighting factor. Assuming a smooth velocity field, we may replace \mathbf{v}_{ij} in equation (B2) with its Taylor expansion $\mathbf{v}_{ij} = \mathbf{V}_i \cdot \mathbf{x}_{ij} + \mathcal{O}(|\mathbf{x}_{ij}|^2)$, where $\mathbf{V}_i \equiv \nabla \otimes \mathbf{v}|_{\mathbf{x}_i}$ is the gradient of \mathbf{v} at position \mathbf{x}_i , and obtain

$$\mathbf{D}_i = \mathbf{V}_i \cdot \mathbf{T}_i + \text{h.o.t.}, \quad (\text{B3})$$

with the symmetric matrix

$$\mathbf{T}_i = \sum_j \mathbf{x}_{ij} \otimes \mathbf{x}_{ij} \tilde{w}_{ij}. \quad (\text{B4})$$

Comparing (B2) and (B4) to the simple estimator (B1), we see that the latter corresponds to (conveniently dropping the index i) $\widehat{\nabla \cdot \mathbf{v}} = v \text{tr}(\mathbf{D}) / \text{tr}(\mathbf{T})$ and the weights $\tilde{w}_{ij} = m_j \tilde{w}_{ij}$. If we split \mathbf{V} into its isotropic part (divergence), the symmetric traceless part \mathbf{S} (shear), and the antisymmetric part \mathbf{R} (vorticity),

$$\mathbf{V} = v^{-1} \nabla \cdot \mathbf{v} \mathbf{I} + \mathbf{S} + \mathbf{R}, \quad (\text{B5})$$

and insert it into (B3), we find for the simple estimator (B1):

$$\widehat{\nabla \cdot \mathbf{v}} = \nabla \cdot \mathbf{v} + v \text{tr}(\mathbf{S} \cdot \tilde{\mathbf{T}}) / \text{tr}(\mathbf{T}) + \text{h.o.t.}, \quad (\text{B6})$$

where $\tilde{\mathbf{T}}$ denotes the anisotropic (traceless) part of \mathbf{T} . Thus, the simple estimator (B1) contains an $\mathcal{O}(h^0)$ error term, which originates from anisotropy of \mathbf{T} in conjunction with velocity shear (owing to the symmetry of \mathbf{T} the vorticity is harmless). For perfectly symmetric particle distributions $\tilde{\mathbf{T}} = 0$, but in general $\tilde{\mathbf{T}} \neq 0$ such that in the presence of strong shear even a small residual $\tilde{\mathbf{T}}$ results in a failure of the simple estimator (B1). This typically happens in differentially rotating discs, where (i) the velocity field is divergent free but contains shear and (ii) even in the absence of noise $\tilde{\mathbf{T}} \neq 0$ owing to the shearing particle distribution.

B2 A more accurate $\nabla \cdot \mathbf{v}$ estimator

From equation (B3), we can also estimate

$$\hat{\mathbf{V}}_i = \mathbf{D}_i \cdot \mathbf{T}_i^{-1}, \quad (\text{B7})$$

which allows an improved divergence estimator

$$\widehat{\nabla \cdot \mathbf{v}}_i = \text{tr}(\mathbf{D}_i \cdot \mathbf{T}_i^{-1}). \quad (\text{B8})$$

In order to assess the error of this estimator, let us expand the flow to second order, replacing equation (B3) with (dropping the index i and using suffix instead of matrix notation)

$$\mathbf{D}_{\alpha\beta} = v_{\alpha,\gamma} \mathbf{T}_{\gamma\beta} - \frac{1}{2} v_{\alpha,\gamma\delta} \mathbf{U}_{\gamma\delta\beta} + \text{h.o.t.}, \quad (\text{B9})$$

with the symmetric tensor $\mathbf{U}_i = \sum_j \mathbf{x}_{ij} \otimes \mathbf{x}_{ij} \otimes \mathbf{x}_{ij} \tilde{w}_{ij}$. Inserting this into (B7) we find

$$\hat{\mathbf{V}}_{\alpha\beta} = v_{\alpha,\beta} - \frac{1}{2} v_{\alpha,\gamma\delta} \mathbf{U}_{\gamma\delta\eta} \mathbf{T}_{\eta\beta}^{-1} + \text{h.o.t.} \quad (\text{B10})$$

Thus, while this estimator avoids an $\mathcal{O}(h^0)$ error, we still have an $\mathcal{O}(h^1)$ error term (since \mathbf{U} is one order higher in h than \mathbf{T}). We can reduce the $\mathcal{O}(h^1)$ error by a careful choice of the weights \tilde{w}_{ij} . If, for instance, $\tilde{w}_{ij} = m_j \tilde{w}_{ij} / \hat{\rho}_i$ then $\mathbf{U} \rightarrow 0$ to leading order in the *continuum limit* by virtue of the isotropy of the kernel. This limit, which is commonly used to assess SPH estimators, replaces $\sum_j m_j$ with $\int \rho(\mathbf{x}_j) d\mathbf{x}_j$ under the assumption of a smooth density without particle noise.⁷ As these conditions are hardly ever truly satisfied, we can only reduce but not eliminate the $\mathcal{O}(h^1)$ error term – as we do not even try to avoid the $\mathcal{O}(h^2)$ error (hidden in ‘h.o.t.’ above), such a reduction should be okay in most cases.

⁷ Under these conditions also $\tilde{\mathbf{T}}$, which causes the $\mathcal{O}(h^0)$ error term in the simple $\nabla \cdot \mathbf{v}$ estimator, vanishes.

B3 Estimating $\hat{\nabla} \cdot \mathbf{v}$

We can estimate $\hat{\nabla} \cdot \mathbf{v}$ either from the change in the estimated $\nabla \cdot \mathbf{v}$ over the last time-step or as the trace of $\hat{\mathbf{V}}$, the total time derivative of \mathbf{V} . Since (with $\mathbf{A} \equiv \nabla \otimes \dot{\mathbf{v}}$ the gradient of the acceleration)

$$\dot{\mathbf{V}} = \mathbf{A} - \mathbf{V}^2 \quad (\text{B11})$$

(a good exercise for your undergraduate students), we can estimate

$$\widehat{\nabla \cdot \mathbf{v}}_i = \text{tr} \left(\hat{\mathbf{A}}_i - \hat{\mathbf{V}}_i^2 \right). \quad (\text{B12})$$

Here, the estimate $\hat{\mathbf{A}}_i$ is obtained from the accelerations at the previous time-step in a way analogous to estimating $\hat{\mathbf{V}}_i$, in particular

we need to compute the matrix \mathbf{T}_i and its inverse only once. The lowest order error in this estimate again is $\mathcal{O}(h^1) \propto \mathbf{U}_i$, such that reducing \mathbf{U}_i by careful choice of the weights remains a good idea.

Note that, by virtue of equation (B11), we could estimate $\hat{\nabla} \cdot \mathbf{v}$ also as $\nabla \cdot \dot{\mathbf{v}} - \text{tr}(\mathbf{V}^2)$ with the acceleration divergence $\nabla \cdot \dot{\mathbf{v}}$ estimated using the standard divergence estimator, in the hope that its $\mathcal{O}(h^0)$ error term is small since the acceleration is hardly sheared.

This paper has been typeset from a \LaTeX file prepared by the author.