

Are heuristic switches necessary to control dissipation in modern smoothed particle hydrodynamics?

Domingo García-Senz^{1,2,*} and Rubén M. Cabezón^{3,4,*}

¹ Departament de Física, Universitat Politècnica de Catalunya (UPC), Av. Eduard Maristany 16, 08019 Barcelona, Spain

² Institut d'Estudis Espacials de Catalunya (IEEC), 08860 Castelldefels (Barcelona), Spain

³ Center for Scientific Computing – sciCORE, University of Basel, Klingelberstrasse 61, 4056 Basel, Switzerland

⁴ Center for Data Analytics – CEDA, University of Basel, 4056 Basel, Switzerland

Received 22 October 2025 / Accepted 24 February 2026

ABSTRACT

Context. Artificial viscosity is commonly employed in smoothed particle hydrodynamics (SPH) to model dissipation in hydrodynamic simulations. However, its practical implementation today relies, in many cases, on complex numerical switches to restrict its application to regions where dissipation is physically warranted, such as shocks. These switches, while essential, are imperfect and can introduce additional numerical noise.

Aims. We investigated an efficient shock capture scheme for SPH that does not rely on artificial viscosity switches. The advantages of the proposed scheme have been validated through a representative number of test cases.

Methods. Recent studies have proposed that subtracting the linear component of the velocity field can suppress spurious dissipation in shear-dominated regions. Building on this idea, we implemented a velocity-reconstruction technique that removes the bulk linear motion from the local velocity field and uses the Balsara correction to modulate the dissipation.

Results. The methodology presented here yields a balanced dissipation scheme that performs well across a range of regimes, including subsonic instabilities, shear flows, and strong shocks. We demonstrate that this approach yields improved accuracy and lower spurious dissipation, compared to the reference viscosity switch used in this work.

Key words. hydrodynamics – turbulence – methods: numerical

1. Introduction

Smoothed particle hydrodynamics (SPH) is a Lagrangian method widely used to simulate fluid flows, originally introduced by [Gingold & Monaghan \(1977\)](#) and [Lucy \(1977\)](#). It is particularly useful in astrophysical contexts where complex geometries and free surfaces are common. Despite its versatility and continued development, SPH still faces challenges in certain areas; notably, in the accurate treatment of dissipation.

Unlike grid-based methods, which typically employ Riemann solvers ([Godunov 1959](#)) to capture fluid discontinuities, SPH relies primarily on artificial viscosity (AV) formulations¹. The theoretical foundations of the standard AV scheme have long been established ([Monaghan & Gingold 1983](#); [Monaghan 1992](#)). However, practical applications of the method still require additional refinements in cases where the required algorithm needs to be sufficiently versatile to handle different fluid regimes.

One of these refinements has been the introduction of corrective factors (AV limiters) that suppress viscosity in strongly shearing regions, thereby reducing dissipation in vorticity-dominated flows and facilitating the growth of hydrodynamic instabilities ([Balsara 1995](#)). A widely adopted and often more effective strategy employs time-dependent dissipation

coefficients (i.e., switches), which remain low by default, rising rapidly when a shock is detected and relaxing back to their nominal values once the shock has passed ([Morris & Monaghan 1997](#); [Monaghan 2005](#)). The different implementations of those time-dependent AV parameters primarily differ in the design and responsiveness of the shock detector and the relaxation law ([Cullen & Dehnen 2010](#); [Read & Hayfield 2012](#); [Rosswog 2015, 2020a](#)).

Although artificial-viscosity switches (AVSWs) represent a major advance and do indeed reduce unwanted dissipation, they are not without drawbacks (however, see [Rosswog 2020a](#), for an entropy-steered variant). In post-shock regions, oscillations in the time-dependent dissipation coefficients can induce excessive numerical noise and spuriously increase entropy. Moreover, for certain problems, such as instability growth and subsonic turbulence, the reported convergence rates are often lower than those achieved with mesh-based methods ([Bauer & Springel 2012](#)).

Recently, [Frontiere et al. \(2017\)](#) proposed the use of slope-limited reconstruction (SLR) within the artificial dissipation terms to suppress spurious viscosity. Conceptually, SLR acts as a local filter that subtracts the linear component of the velocity field, preventing bulk shear from feeding the dissipative operator. The resulting algorithm is not derived from first principles and is therefore heuristic in nature as well. However, unlike AVSW, it is purely local and instantaneous: it involves no time derivatives and does not require prescribing a working range for the viscosity coefficient α . In previous studies, SLR has been combined with AVSW ([Frontiere et al. 2017](#); [Rosswog 2020a](#); [Rosswog & Diener 2025](#); [Cabezón et al. 2025](#)) with good effect.

* Corresponding authors: domingo.garcia@senior.upc.edu;
ruben.cabezon@unibas.ch

¹ There have also been proposals to adapt Riemann solvers for use in SPH ([Inutsuka 2002](#); [Monaghan 1997](#); [Cha & Whitworth 2003](#); [Cha et al. 2010](#)), representing a promising parallel line of development ([Rosswog 2024](#)).

Nevertheless, several tests indicate that further gains are achieved when AVSWs are omitted altogether (Sandnes et al. 2025). Similarly, Price (2024) explored switch-free dissipation and reported that SLR compares favorably to switch-based schemes, albeit only for 1D test problems. The prospect of attaining high accuracy without AV switches is noteworthy and given the ubiquity of these terms in SPH, it warrants closer examination.

On another note, SLR schemes and, in general, algorithms based on total variation diminishing (TVD) approaches have been criticized for not being automatically consistent with the entropy inequality (Yee et al. 1983). In SLR schemes, dissipation is not guaranteed to be positive-definite by construction (Price & Laibe 2020), which may cause local violations of the second principle of thermodynamics. There are ways to modify TVD to be consistent with the entropy condition in 1D (Harten et al. 1976; Osher & Chakravarthy 1984), but this remains an open question in 3D. However, as discussed in Sect. 5, we have not seen anomalous entropy behavior in any of the tests performed when SLR was used². Other situations where switches can potentially outperform Godunov-SPH and SLR methods (specifically, 1D linear-wave evolution) have been discussed in Puri & Ramachandran (2014a).

In this work, we carried out a suite of SPH simulations using a SPH formulation distinct from that of Sandnes et al. (2025), with two goals: (i) to independently validate the performance of the slope-limited reconstruction (SLR) technique; and (ii) to refine it where it can prove beneficial. Our test suite included several benchmarks considered by Sandnes et al. (2025) as well as additional problems that probe shear- and turbulence-dominated regimes. We have paid particular attention to turbulent flows, where excessive artificial dissipation is especially detrimental.

This article is organized as follows. Section 2 summarizes the different dissipation schemes currently used to handle dissipation as well as our proposed methods. Section 3 presents various tests comparing the performance of these schemes. Section 4 is entirely dedicated to the study of subsonic turbulence and how the dissipation schemes influence SPH's ability to resolve it. The methods that offer the best balanced results across all tests are ranked in Sect. 5, where the potential disadvantages of SLR schemes are also outlined and discussed. Finally, in Sect. 6, we discuss our main findings, along with the recommended shock-capturing setting and future prospects.

2. SPH flavor

Most of the test cases described in Sect. 3 were carried out with the SPH code SPHYNX (Cabezón et al. 2017). Only the turbulence simulations were performed with SPH-EXA (Cavelan et al. 2020; Keller et al. 2023), due to its scalability and native GPU acceleration. This enabled higher-resolution, time- and resource-intensive runs to be executed substantially faster. Both codes incorporate the same improved algorithms for gradient estimation (García-Senz et al. 2012; Cabezón et al. 2012; Rosswog 2015), kernel interpolation (Cabezón et al. 2008; Cabezón & García-Senz 2024), and an accurate partition of unity (García-Senz et al. 2022), while the equation of energy is integrated consistently to describe thermal evolution (i.e., grad-h terms are included). These ingredients are fully consistent with the

² To be precise, both the SLR and the AV-switches showed slight entropy inversions in noise-dominated regions and around contact discontinuities when these were not conveniently smoothed (e.g., in the shock-tube test).

Lagrangian nature of SPH, yielding a conservative scheme. The remainder of this section focuses on the construction and implementation of dissipation in artificial viscosity-based SPH schemes.

2.1. Artificial viscosity schemes

As in other modern SPH formulations (Price et al. 2018; Rosswog 2020b; Wadsley et al. 2017), SPHYNX and SPH-EXA handle dissipation with the classical artificial viscosity (AV) scheme (Monaghan 1992, 1997). In particular, Monaghan (1997) defines the pairwise AV term Π_{ab} in analogy to a Riemann problem,

$$\Pi_{ab} = -\frac{\alpha v_{ab}^{sig} \omega_{ab}}{2 \bar{\rho}_{ab}}, \quad (1)$$

with

$$\omega_{ab} = \mathbf{v}_{ab} \cdot \hat{\mathbf{r}}_{ab}, \quad (2)$$

where $\mathbf{v}_{ab} = \mathbf{v}_a - \mathbf{v}_b$ and $\hat{\mathbf{r}}_{ab}$ is the unit vector joining particles a and b . Here,

$$v_{ab}^{sig} = c_a + c_b - \gamma \omega_{ab} \quad (3)$$

is a characteristic signal velocity at which the information propagates between particles a and b , and $\gamma = 3$. Equations (1) and (3) can be loosely related to the widely used expression of Monaghan (1992), namely,

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

where

$$\mu_{ab} = \frac{\bar{h}_{ab} \mathbf{v}_{ab} \cdot \mathbf{r}_{ab}}{\mathbf{r}_{ab}^2 + \eta^2}. \quad (5)$$

Neglecting the regularization, η , and taking $\frac{\bar{h}_{ab}}{|\mathbf{r}_{ab}|} = 1$ yields $\mu_{ab} \approx \omega_{ab}$; thus, comparing Eqs. (1) and (4) gives $\beta = \gamma \alpha / 2$. For $\gamma = 3$, we have $\beta = \frac{3}{2} \alpha$. However, a more frequent choice is $\beta = 2\alpha$, which better reduces post-shock oscillations in strong shocks. In the continuum limit, it has been shown that the linear part of AV in Eq. (4) is equivalent to the shear and bulk viscosities in the Navier-Stokes equations (Monaghan 2005). Thus, the AV is not as ‘‘artificial’’ as it might seem after all.

Modern schemes handle α and β independently, since α is often allowed to vary per particle and in time during calculations, while β is kept constant. For example, Wadsley et al. (2017) considered

$$\Pi_{ab} = \begin{cases} \frac{-\bar{\alpha}_{ab} \bar{c}_{ab} \omega_{ab} + \beta \omega_{ab}^2}{\bar{\rho}_{ab}} & \text{for } \omega_{ab} < 0, \\ 0 & \text{for } \omega_{ab} \geq 0, \end{cases} \quad (6)$$

where ω_{ab} is given by Eq. (2) and $\alpha_{min} \leq \alpha_a \leq \alpha_{max}$, while β remains constant. SPHYNX and SPH-EXA also employ Eq. (6) to implement dissipation. Typical values of the AV coefficients when using switches are: $\alpha_{min} = 0.05$, $\alpha_{max} = 1$, and $\beta = 2$.

2.2. Artificial viscosity with shock detection: Switches

The choice of α and β , which control the linear ($\propto \omega_{ab}$) and the quadratic ($\propto \omega_{ab}^2$) contributions to the AV in Eq. (6), is crucial in practical applications of SPH. The simplest option is to keep α and β constant during the calculation. However, values of α large enough to capture strong shocks introduce excessive viscosity in smooth regions of the fluid. Overdissipation damps short-wavelength modes and hinders the growth of fluid instabilities and turbulence. It also suppresses differential rotation, resulting in a poor depiction of, for instance, the accretion disk dynamics.

A plausible, albeit phenomenological, remedy was proposed by Balsara (1995): adopting constant $\alpha = 1$ and $\beta = 2$, but multiplying Π_{ab} by a limiter \mathcal{B}_{ab} designed to approach unity in compressive (shock) regions and to decrease toward zero in curl-dominated flows,

$$\Pi_{ab} \leftarrow \mathcal{B}_{ab} \Pi_{ab}; \quad \mathcal{B}_{ab} = \frac{1}{2}(\mathcal{B}_a + \mathcal{B}_b), \quad (7)$$

with

$$\mathcal{B}_a = \frac{|\nabla \cdot \mathbf{v}|_a}{|\nabla \cdot \mathbf{v}|_a + |\nabla \times \mathbf{v}|_a + 10^{-4} \frac{c_a}{h_a}}, \quad (8)$$

where c_a and h_a are the speed of sound and the smoothing length of particle a , respectively.

Later, a more efficient algorithm was introduced by Morris & Monaghan (1997) and subsequently refined by others. In this approach, the linear AV coefficient is kept near a low default value, $\alpha = \alpha_{min}$ (with $\alpha_{min} \simeq 0.05$), but quickly increases towards α_{max} when a shock is detected. Cullen & Dehnen (2010) showed that fast and robust shock detection is crucial in this technique and proposed using a higher order indicator based on the derivative of $\nabla \cdot \mathbf{v}_a$, which generally yields improved results. Several variations and refinements of these switches for controlling $\alpha(\mathbf{r}, t)$ can be found in the literature (Read & Hayfield 2012; Rosswog 2015). By default, SPHYNX and SPH-EXA use the same high-order dissipation switch proposed in Read & Hayfield (2012), which can detect flow convergence before it occurs and can be considered the state-of-the-art version of that technique.

The specific implementation is the following: α_a is initialized at a low value (usually $\alpha_a = 0.05$). Then it increases rapidly if dissipation is needed and decays exponentially to its original value when it is not needed, effectively acting as a switch for the AV at the position of each SPH particle.

We start by calculating a local viscous parameter, $\alpha_{loc,a}$ (Read & Hayfield 2012), expressed as

$$\alpha_{loc,a} = \begin{cases} \frac{\alpha_{max} \mathcal{D}_a}{\mathcal{D}_a + h_a |\nabla \cdot \mathbf{v}_a| + 0.05 c_a} & \text{for } \nabla \cdot \mathbf{v}_a < 0, \\ 0 & \text{otherwise,} \end{cases} \quad (9)$$

where $\mathcal{D}_a = h_a^2 |\nabla(\nabla \cdot \mathbf{v}_a)|$ predicts flow convergence by comparing the gradient of $\nabla \cdot \mathbf{v}_a$ with the local value of $\nabla \cdot \mathbf{v}_a$. If the former is large enough and the fluid is actually converging ($\nabla \cdot \mathbf{v}_a < 0$), $\alpha_{loc,a}$ adopts a value between 0 and α_{max} . Furthermore, the term $0.05 c_a$, where c_a stands for the local speed of sound, acts as a base value for the size of the velocity fluctuations that can trigger dissipation, and it also works as a term to prevent divergent $\alpha_{loc,a}$. Once we know $\alpha_{loc,a}$ we have two options:

- if $\alpha_{loc,a} \geq \alpha_a$: we instantaneously raise the dissipation setting $\alpha_a = \alpha_{loc,a}$,
- if $\alpha_{loc,a} < \alpha_a$: we let α_a decay smoothly.

In the case that α_a has to decay, we evaluate an instantaneous decay rate of $\dot{\alpha}_a$ (Read & Hayfield 2012), expressed as

$$\dot{\alpha}_a = \begin{cases} (\alpha_{loc,a} - \alpha_a)/\tau_a & \text{if } \alpha_{min} < \alpha_{loc,a} < \alpha_a, \\ (\alpha_{min} - \alpha_a)/\tau_a & \text{if } \alpha_{min} \geq \alpha_{loc,a}. \end{cases} \quad (10)$$

The decay time, τ_a , is defined as $\tau_a = h_a/v_{sig,max}$ and it is a function of the local spatial resolution and the maximum signal velocity, $v_{sig,max}$, in the neighborhood of particle a , calculated according to Eq. (3). The effect of Eq. (10) is to allow alpha to decay exponentially to $\alpha_{loc,a}$ (or α_{min}) if $\alpha_{loc,a}$ is too low. In this way, we can ensure a minimum dissipation that helps reduce the random noise that might arise from particle disorder. However, we note that even though all particles have at least a minimum value, α_{min} , dissipation is automatically cut off when $\mathbf{r}_{ab} \cdot \mathbf{v}_{ab} \geq 0$. Also, there is no restriction for α_{min} , which can be set to zero.

At this point, it is worth highlighting the subtle difference between a limiter and a switch in this context. A limiter, such as the Balsara one, is a continuous, instantaneous, memoryless amplitude control that keeps dissipation high by default, except in regions that meet well-defined physical criteria (e.g., a high curl value). A switch, on the other hand, is a tunable sensor with temporary memory that can lag or overshoot in different regions. Switches are engineered to work well in some tests, but their exact form is a practical choice, rather than a unique consequence of the equations.

Despite their success, AVSWs can introduce additional noise, especially in the post-shock wake where α may oscillate due to rapid changes in the shock indicator, leading to spurious entropy production (e.g., Cullen & Dehnen 2010; Read & Hayfield 2012; Rosswog 2015). In mixed regimes where shocks and strong shear coexist (e.g., when a dilute stream impacts a dense cloud), switches might struggle to balance compression-driven dissipation against shear suppression (cf. Balsara 1995; Price et al. 2018). Finally, switches tend to be less effective in controlling low-level numerical noise in otherwise quiescent regions; dedicated “noise triggers” can be added to mitigate this, but at the expense of additional heuristics and computing time (Rosswog 2015).

2.3. A shock-capturing algorithm free of switches

A linear velocity field implies a smooth flow without the presence of shocks. Therefore, linear velocity fields should have a vanishing AV. Following the prescription introduced by Frontiere et al. (2017), we reconstructed the local velocity field for each particle by removing the linear component of the velocity jump that is used to trigger and calculate the dissipation (Eq. (6)). Assuming differentiability in the neighborhood of each SPH particle, the optimal linear approximation of the velocity field is given by its Jacobian acting on the separation vector. Rather than using $\mathbf{v}_{ab} = \mathbf{v}_a - \mathbf{v}_b$ in the AV, we use $\mathbf{v}'_{ab} = \mathbf{v}'_a - \mathbf{v}'_b$, defined as

$$\mathbf{v}'_{a,i} \equiv v_{a,i} - \frac{1}{2} \phi_{ab} \mathbf{J}_{\mathbf{v}_a} \mathbf{x}_{ab}^T, \quad (11)$$

$$\mathbf{v}'_{b,i} \equiv v_{b,i} + \frac{1}{2} \phi_{ba} \mathbf{J}_{\mathbf{v}_b} \mathbf{x}_{ab}^T, \quad (12)$$

where $\mathbf{x}_{ab}^T = \mathbf{x}_a - \mathbf{x}_b$ is the column vector of the relative distance between particles a and b , and $\mathbf{J}_{\mathbf{v}}$ is the Jacobian of the velocity field. In the presence of discontinuities, such a linear estimator is not accurate enough; hence, the factor ϕ_{ab} is designed to approach zero in these cases and to remain close to unity for

Table 1. AV variants considered in this work.

Method	SLR	\mathcal{B} modulation	Switches	α coefficient	Notes
1: AV	No	–	No	$\alpha = 1$	Classical shock-capturing baseline.
2: AVSW	No	–	Yes	$\alpha_a \in [\alpha_{\min}, \alpha_{\max}]$	Time-dependent α .
3: AVSLR	Yes	No	No	$\alpha = 1$	Slope-limited reconstruction.
4: AVSWSLR	Yes	No	Yes	$\alpha_a \in [\alpha_{\min}, \alpha_{\max}]$	Hybrid: switches + SLR.
5: AVSLRB	Yes	$(1 - \mathcal{B})$	No	$\alpha = 1$	No SLR when $\mathcal{B} \simeq 1$; strong SLR when $\mathcal{B} \simeq 0$.
6: AVSLRB2	Yes	$(1 - \mathcal{B}^2)$	No	$\alpha = 1$	No SLR when $\mathcal{B} \simeq 1$; strong SLR if $\mathcal{B} \lesssim 0.3$.
7: AVSLRB2 \mathcal{F}	Yes	$(1 - \mathcal{B}^2)$	No	$\alpha = \alpha_{\max} \times \max[\mathcal{F}, \mathcal{B}]$	Hybrid: SLR + α with floor value \mathcal{F} .

Notes. SLR denotes the slope-limited velocity reconstruction of Eqs. (11)–(12), optionally modulated by the Balsara factor \mathcal{B} ($0 \leq \mathcal{B} \leq 1$) via Eqs. (18)–(19). Switches refers to time-dependent $\alpha(\mathbf{r}, t)$. The last row is discussed in relation to the subsonic turbulence scenario (Sect. 4). The coefficient β is always fixed to $\beta = 2$.

smooth flow. [Frontiere et al. \(2017\)](#) proposed a van Leer-like limiter ([van Leer 1974](#)):

$$\phi_{ab} = \max \left[0, \min \left(1, \frac{4F_{ab}}{(1 + F_{ab})^2} \right) \right] \times \kappa_{ab}, \quad (13)$$

$$\kappa_{ab} \equiv \begin{cases} \exp \left\{ - \left(\frac{q_{ab} - q_{crit}}{q_{fold}} \right)^2 \right\} & \text{if } q_{ab} < q_{crit}, \\ 1 & \text{if } q_{ab} \geq q_{crit}, \end{cases} \quad (14)$$

$$q_{ab} \equiv \min(q_a, q_b), \quad (15)$$

$$q_{crit} \equiv \left(\frac{32\pi}{3nb_a} \right)^{1/3}, \quad (16)$$

$$F_{ab} \equiv \frac{\mathbf{x}_{ab} \mathbf{J}_{\mathbf{v}_a} \mathbf{x}_{ab}^T}{\mathbf{x}_{ab} \mathbf{J}_{\mathbf{v}_b} \mathbf{x}_{ab}^T}, \quad (17)$$

where $q_a = |\mathbf{x}_b - \mathbf{x}_a|/h_a$, q_{crit} estimates the mean interparticle spacing in units of smoothing length, while nb_a is the neighbor count of particle a . The parameter q_{fold} sets the width of the Gaussian correction (we adopted the recommended value from [Frontiere et al. \(2017\)](#) of $q_{fold} = 0.2$).

Equations (11) and (12) remove the linear component of the velocity field, limiting the slope at discontinuities, while reducing unwanted dissipation in smooth flows. We refer to this slope-limited reconstruction as AVSLR, and contrast it with our implemented AVSW (time-dependent switch) method in what follows. We note that although both AVSLR and AVSW share a heuristic origin, linear reconstruction does not involve time derivatives and decay laws.

Currently, few calculations incorporating these AVSLRs have been reported ([Frontiere et al. 2017](#); [Rosswog 2020a,b](#); [Cabezón et al. 2025](#); [Sandnes et al. 2025](#)). They were mostly used in combination with switches. As in [Price \(2024\)](#) and [Sandnes et al. \(2025\)](#), our premise is that AV switches are unnecessary when using the slope-limited reconstruction algorithm. We therefore explored several variants of the implementation of the SLR (see Table 1), as follows:

- AVSLR: use Eqs. (11) and (12) to remove the linear component of the velocity field;
- AVSLRB: modulate the SLR with the Balsara limiter:

$$v'_{a,i} \equiv v_{a,i} - \frac{1}{2} (1 - \mathcal{B}_a^p) \phi_{ab} \mathbf{J}_{\mathbf{v}_a} \mathbf{x}_{ab}^T, \quad (18)$$

$$v'_{b,i} \equiv v_{b,i} + \frac{1}{2} (1 - \mathcal{B}_b^p) \phi_{ba} \mathbf{J}_{\mathbf{v}_b} \mathbf{x}_{ab}^T, \quad (19)$$

- with $\mathcal{B}_{a,b}$ defined by Eq. (8)³. Following the suggestion of [Sandnes et al. \(2025\)](#), we first consider the case with $p = 1$.
- AVSLRB2 same as above but with $p = 2$; this does not change the description of shocks, although it proved more advantageous than $p = 1$ in our tests.
- Baselines: for comparison, we also consider (i) classical dissipation with fixed coefficients (AV); (ii) AV with time-dependent switches (AVSW); and (iii) a hybrid combining switches with slope-limiting (AVSWSLR). The objective is to find the most balanced and general scheme among all these options in light of the results.

3. Tests

We first considered six 3D test problems designed to compare the performance of the proposed schemes across shock- and shear-dominated regimes. The shock-dominated set comprises the Sod shock tube ([Sod 1978](#)) and the Sedov blast ([Sedov 1959](#)). The shear and instability set comprises the steady Gresho–Chan vortex ([Gresho & Chan 1990](#)), Kelvin–Helmholtz (KH), and Rayleigh–Taylor (RT) instabilities, as well as the wind–cloud interaction where a dilute supersonic wind disrupts a dense cloud ([Agertz et al. 2007](#)). A seventh test dealing with subsonic turbulence is deferred to Sect. 4. When needed, we adopted an ideal-gas equation of state $P = (\gamma - 1)\rho u$, with $\gamma = 5/3$ (or with $\gamma = 1.4$ and $\gamma = 1.001$ for the RT and turbulence tests, respectively). Simulations were carried out with SPHYNX⁴ ([Cabezón et al. 2017](#)), with the exception of turbulence, which was calculated with SPH-EXA⁵, (including the updates discussed in [García-Senz et al. 2022](#)). Kernel interpolations use the mixed sinc kernel $S_{4,9}^{0.9}$ ([Cabezón & García-Senz 2024](#)), which is pairing-resistant and robust for a wide range of neighbors, $60 \leq n_b \leq 400$. Unless stated otherwise, we set $n_b = 100$. Additional setup details for each case are summarized in Tables 1 and 2.

To make quantitative comparisons between models, we report a mean absolute error L_1 wherever a suitable reference is available,

$$L_1 = \frac{1}{N_S} \sum_{i \in S} |p_i^{\text{SPH}} - p_i^{\text{ref}}|, \quad (20)$$

³ In the tests performed here, we used the average $\mathcal{B}_{ab} = \frac{1}{2}(\mathcal{B}_a + \mathcal{B}_b)$ in Eqs. (18) and (19) so that the multiplicative factor $\frac{1}{2}(1 - \mathcal{B}_{ab}^p)$ is the same in both expressions.

⁴ <https://astro.physik.unibas.ch/sphynx>

⁵ <https://github.com/sphexa-org/sphexa>

Table 2. Simulations name convention and some features of the studied test cases.

Test	AV $\alpha = 1$	AVSW $\alpha(t)$	AVSLR $\alpha = 1$	AVSWSLR $\alpha(t)$	AVSLRB $\alpha = 1$	AVSLRB2 $\alpha = 1$	AVSLRB2 \mathcal{F} $\alpha = 1, \mathcal{F} = 0.5$	Box $[L_x, L_y, L_z]$	N ($\times 10^6$)	N^{1D}
Shock-Tube	ST ₁	ST ₂	ST ₃	ST ₄	ST ₅	ST ₆	ST ₇	[1, 1/8, 1/8]	0.97	396
Sedov	S ₁	S ₂	S ₃	S ₄	S ₅	S ₆	S ₇	[1, 1, 1]	8.0	200
GH-Vortex	V ₁	V ₂	V ₃	V ₄	V ₅	V ₆	V ₇	[1, 1, 0.5]	0.5	100
KHI	KH ₁	KH ₂	KH ₃	KH ₄	KH ₅	KH ₆	KH ₇	[1, 1, 0.1]	0.6	182
RTI	RT ₁	RT ₂	RT ₃	RT ₄	RT ₅	RT ₆	RT ₇	[1, 2, 0.1]	1.2	182
Wind-Cloud	WC ₁	WC ₂	WC ₃	WC ₄	WC ₅	WC ₆	WC ₇	[40, 10, 10]	0.5	200
Turbulence	T ₁	T ₂	T ₃	T ₄	T ₅	T ₆	T ₇	[1, 1, 1]	512	800

Notes. The first column is the test type. The second to seventh columns show the test name that relates the test type to the methods defined in Table 1. The three last columns indicate the size of the computational box, the total number of SPH particles, and the equivalent number per axis, N^{1D} in a cubic box with $L = 1$, calculated as $N^{1D} = [N/(L_x \cdot L_y \cdot L_z)]^{1/3}$. $\beta = 2$ in all tests.

where S denotes the set of sampling points with positions d_i in the range $d_0 \leq d_i \leq d_f$ (e.g., a 1D line-out or radial profile), and N_S is the number of particles belonging to the set, S . Here, p^{SPH} is the measured quantity from the simulation and p^{ref} is the corresponding reference (analytic solution or high-resolution benchmark) specified for each test below.

3.1. The shock-tube

In this test, a box with the size and number of equal-mass particles specified in Table 2 is initialized with a gas that has a strong discontinuity in density and pressure across a contact interface at ($x = x_c$). The problem has a well-known analytical solution (Sod 1978), which we use to assess the performance of the different shock-capturing prescriptions.

The setup follows common practice (Wadsley et al. 2017; Price et al. 2018): in the left state ($x < x_c$), we take $\rho_l = 1$, $p_l = 1$; in the right state ($x > x_c$), $\rho_r = 0.125$, $p_r = 0.10$. The particle distribution across the contact is left unsmoothed, but the pressure was smoothed according to

$$P_0 = -A \tanh\left(\frac{x - x_c}{\delta}\right) + B, \quad (21)$$

with $\delta = 0.01$, $x_c = 0.5$, $A = 0.45$, and $B = 0.55$.

This test probes how the different AV choices handle strong shocks. The results are summarized in Fig. 1, which shows the x -velocity profile, v_x . We computed the L_1 error using Eq. (20) over $x \in [x_0, x_f]$ with $x_0 = d_0 = 0.05$ and $x_f = d_f = 0.868$. The upper limit of $x_f = 0.868$ was chosen to exclude the terminal vertical drop in the velocity, so that the sum focuses on the resolved shock, contact, and rarefaction features.

From Fig. 1, models ST₂ and ST₄ (i.e., those employing switches) show the worst performance, exhibiting pronounced velocity oscillations in the post-shock region. Substantially better results are obtained either with the simple, constant-parameter AV (ST₁) or with the SLR approach, modulated with the Balsara limiter (ST₅ and ST₆). The ST₃ calculation is intermediate, with oscillations confined to the immediate shock vicinity. Overall, the best performance in this set is delivered by ST₁, followed by ST₆, which uses SLR and applies the squared ($p = 2$) Balsara limiter to the functions in Eqs. (18) and (19).

3.2. Spherical explosion

Adopting the test that is typically referred to as the Sedov–Taylor blastwave, we performed it in a 3D box with a side $L = 1$ with

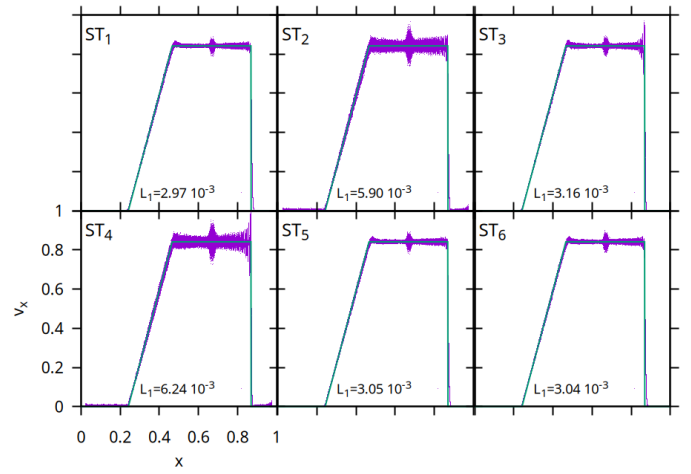


Fig. 1. Velocity profile, v_x , for the shock-tube test and the models shown in the second row in Table 2, at the time $t = 0.2$.

$N = 200^3$ particles arranged in a glass-like configuration. All particles have the same mass, resulting in a homogeneous initial density profile of $\rho = 1 \pm 0.005$. The explosion is initiated in the center of the box by depositing a total energy $\Delta U = 1$ into the gas, distributed with a Gaussian profile of characteristic width $\delta = 2h$, where h is the smoothing length at $t = 0$.

Figures 2 and 3 summarize the results. The time evolution of the maximum density is shown in Fig. 2. According to these figures, models S₃ and S₄ overshoot the strong-shock compression limit $\rho/\rho_0 = (\gamma + 1)/(\gamma - 1) = 4$ for $\gamma = 5/3$, while the remaining four models stay below and approach this limit from below as the blast develops. As in the shock-tube test, we computed the L_1 error. In this case, for the radial velocity between the origin and the r coordinate of the maximum velocity. The lower panels of Fig. 3 show the profile of v_{rad} and the L_1 value at $t = 0.1$. In the central region ($r \leq 0.25$), the numerical profiles depart from the analytical solution (solid line), although the amount of mass involved is very small. To adequately resolve the diluted central region, a much larger number of SPH particles would be needed. In the post-shock region, all the models exhibit velocity oscillations, which are particularly pronounced near the shock front for models S₃ and S₄, consistent with the behavior observed in the shock-tube test.

According to these results, the best models are S₆, S₁, and S₅ (these three have very similar results), followed by S₂. The




Fig. 2. Maximum density evolution in the Sedov test for the range of models in Table 2.

worst results arise when using switches together with the SLR technique (S_4) or with SLR alone without Balsara's modulation (S_3).

3.3. The Gresho–Chan vortex

The simulation of a pure fluid vortex (Gresho & Chan 1990) is challenging for SPH schemes working with standard AV: excessive viscosity damps the differential rotation and degrades the velocity field. The flow is axisymmetric and defined by

$$v_\varphi(r) = v_0 \begin{cases} \psi & \text{for } \psi \leq 1, \\ 2 - \psi & \text{for } 1 < \psi \leq 2, \\ 0 & \text{for } \psi > 2, \end{cases} \quad (22)$$

$$P(r) = P_0 + 4v_0^2 \begin{cases} \psi^2/8 & \text{for } \psi \leq 1, \\ (\psi^2/8 - \psi + \ln \psi + 1) & \text{for } 1 < \psi \leq 2, \\ (\ln 2 - 1/2) & \text{for } \psi > 2, \end{cases} \quad (23)$$

with $\psi = 5r$, $r = \sqrt{x^2 + y^2}$, $v_0 = 1$, $\rho_0 = 1$, and $P_0 = 5$.

The box size and number of SPH particles are listed in Table 2. Because the initial azimuthal velocity is piecewise linear, this test is expected to favor AVSLRs methods over AVSWs. The results are summarized in Fig. 4, which also provides the L_1 error for quantitative comparisons. We compute L_1 using Eq. (20) over $r \in [d_0, d_f] = [0, 0.5]$, with p^{SPH} , p^{ref} taken as the azimuthal velocity from the SPH calculation and the analytic v_φ in Eq. (22), respectively.

Figure 4 clearly shows the superior performance of AVSLRs schemes over AVSWs. As expected, the worst case is V_1 , which uses AV with constant coefficients $\alpha = 1$, $\beta = 2$, and no limiter. Introducing switches yields a substantial improvement, but still falls short of the SLR methods. For this test, the best results are obtained with V_4 (AVSRL), V_5 (AVSLRB), and V_6 (AVSLRB2), with practically no differences between them.

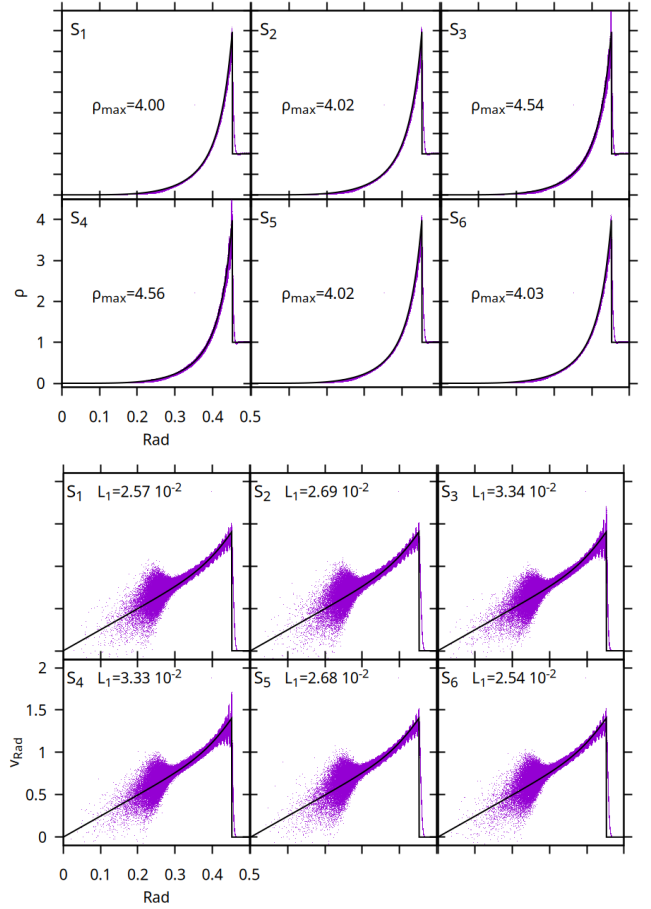


Fig. 3. Density and radial velocity profiles in the Sedov-Taylor test at time $t = 1$ for the different models described in Table 2. The solid line represents the analytical reference. The L_1 error between the numerical and analytical estimates from the center until the maximum radial velocity is also shown.




Fig. 4. Azimuthal velocity distribution in the Gresho–Chan vortex experiment at $t = 1$ for the model V_k in the third row in Table 2. The analytical reference is indicated by the continuum line.

3.4. The Kelvin–Helmholtz instability

The evolution of a shear layer between two fluids of different densities naturally triggers the Kelvin-Helmholtz instability (KHI). Because this instability is ubiquitous in nature, including astrophysical phenomena, its accurate simulation is of great




Fig. 5. Color-coded density for the KH models in Table 2 at $t = 3$ ($t/t_{\text{KH}} = 2.8$).

Table 3. KH: perturbation amplitude $A(t)$ at $t = 1.5$ for the models in Table 2.

–	KH ₁	KH ₂	KH ₃	KH ₄	KH ₅	KH ₆
$A (\times 10^{-2})$	2.87	8.31	12.58	12.14	12.01	12.52

Notes. The reference value from McNally et al. (2012) is $A(t = 1.5) = 14.79 \times 10^{-2}$.

importance. Since AV can artificially dampen KHI growth, this problem is a sensitive test for the SLR algorithm.

We initialize SPH particles in a thin 3D box of the size listed in Table 2, with a 2:1 particle number ratio between the high- and low-density regions. The initial setting consists of three stratified layers, with the central layer having the higher density.

The initial number of neighbors is $n_b = 120$. The initial velocity distribution follows McNally et al. (2012), where shear was induced by imposing $v_x = \pm 0.5$ across the separation layer, exponentially damped along the Y -direction over a characteristic length, $L = 0.025$. The density across the contact layer is not smoothed. We seed a single-mode normal perturbation in the cross-stream velocity,

$$v_y(t = 0) = \Delta v_y \sin\left(\frac{2\pi}{\lambda} x\right), \quad (24)$$

where $\Delta v_y = 0.01$; $\lambda = 0.5$ and we adopted an initial background pressure of $P_0 = 2.5$.

Times are also reported normalized to the characteristic KHI growth time t_{KH} , as defined in Agertz et al. (2007), resulting in $t_{\text{KH}} = 1.06$ for our models. The results are shown in Figs. 5 and 6, and Table 3. Figure 5 shows the density color maps of the different models at a time, $t = 3$ ($t/t_{\text{KH}} = 2.8$). At that time, the instability is fully nonlinear, with well-developed swirls, billows, and secondary structures. These extend over a wide range of wavelengths, with a cut-off set by the real viscosity but, in practice, by the resolution and numerical dissipation.

According to Fig. 5, models KH₃ to KH₆ develop secondary structure more robustly than KH₁ (AV with constant coefficients and no switches). Steering dissipation with switches (KH₂) is the second-worst. This is quantified in Fig. 6, which displays the evolution of the perturbation amplitude. Then, KH₃ and KH₆ show the fastest growth. At $t = 1.5$, the amplitudes (presented in Table 3) are close to the reference (McNally et al. 2012), which




Fig. 6. Amplitude mode growth of the KH models. A zoom-in of the nearby lines KH₃, KH₄, KH₅, and KH₆ around $t = 1.5$ is shown in the subfigure.

is remarkable given the moderate resolution in these simulations $N^{1D} = 182$.

We also see that the use of switches (KH₂), represents a considerable improvement over AV with constant coefficients (KH₁). Even so, the computed growth rate is much lower than that obtained with the SLR family, especially KH₃ and KH₆. The KH₅ and KH₄ models show slightly less amplitude growth than KH₃, but still larger than KH₂, indicating that they can be considered good options when we are looking to balance accuracy and robustness.

3.5. The Rayleigh–Taylor instability

The study of convective-like instabilities in a gravitational field is of major importance in astrophysics. Historically, SPH has struggled to model these systems. In particular, the growth of the Rayleigh–Taylor instability (RTI) was often poorly depicted because AV introduces excessive damping and kills the smallest perturbation modes. However, combining improved gradient estimators with AVSW made a reasonable approach to this phenomenon possible (Cabezón et al. 2017; Rosswog 2020a).

In this test, we used a box of size $[1, 2, 0.1]$ filled with a two-density fluid. The upper half has $\rho_u = 2$, whereas the lower half $\rho_d = 1$. The system is immersed in a constant gravitational field with $g = -0.5$, and the initial pressure profile is set so that the hydrostatic balance is satisfied. The interface between the low- and high-density regions is sharp (no smoothing). The pressure at the bottom boundary is $P(y = 0) = 2.5$. The pressure profile is

$$P(y) = \begin{cases} P(y = 0) + \rho_d g y & y \leq y_b - \frac{\delta}{2}, \\ P\left(y_b - \frac{\delta}{2}\right) - \bar{\rho} g \left(y_b - \frac{\delta}{2}\right) + \bar{\rho} g y & y_b - \frac{\delta}{2} < y < y_b + \frac{\delta}{2}, \\ P\left(y_b + \frac{\delta}{2}\right) + \rho_u g \left(y - \left(y_b + \frac{\delta}{2}\right)\right) & y \geq y_b + \frac{\delta}{2}, \end{cases} \quad (25)$$

where y_b is the interface location, $\bar{\rho}$ is the average density around y_b , and $\delta = 0.4$ is the initial width of the unstable layer $[y_b - \delta/2,$

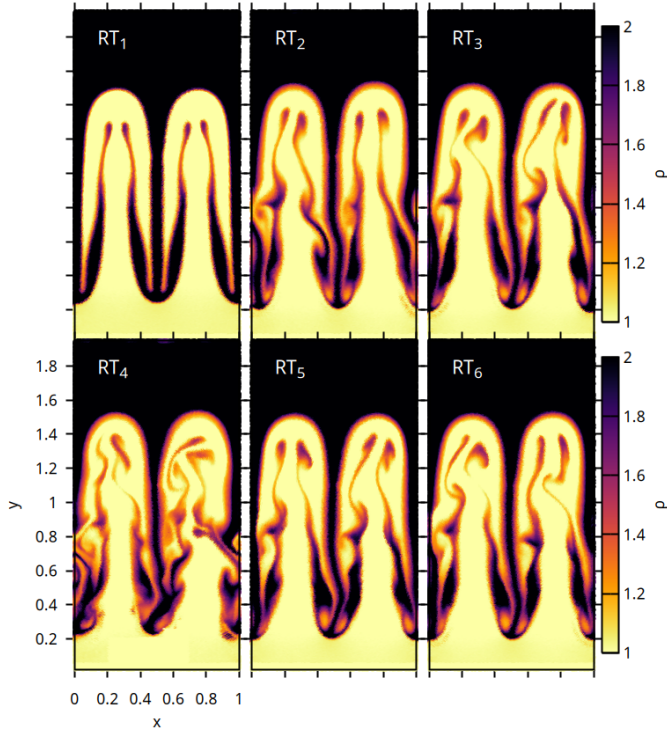


Fig. 7. Color maps of density for the RT models in Table 2 at $t = 6$.

Table 4. RTI: kinetic energy (KE) in the Y -direction at $t = 6$ for the models in Table 2.

–	RT ₁	RT ₂	RT ₃	RT ₄	RT ₅	RT ₆
KE ($\times 10^{-3}$)	6.81	8.49	8.87	8.39	8.62	8.87

$y_b + \delta/2$]. The fluid is initially at rest, and we seed a velocity perturbation, Δv_y , localized to the interface,

$$v_y = \Delta v_y [1 - \cos(2\pi x)] \left[1 + \cos\left(\frac{2\pi(y - y_b)}{\delta}\right) \right], \text{ for } |y - y_b| \leq \frac{\delta}{2}, \quad (26)$$

with $\Delta v_y = 0.025$.

Despite the moderate resolution $N^{1D} = 182$, we can assess the performance by inspecting the density maps in the nonlinear regime and the kinetic energy in the vertical direction. Figure 7 shows the color-coded density at $t = 6$, while Table 4 provides quantitative information on the kinetic energy stored in vertical displacements at the same time. The RTI grows and becomes nonlinear even when with AV using constant $\alpha = 1, \beta = 2$ (RT₁), but without any discernible substructure. Introducing a time-dependent α (RT₂) significantly improves the outcome, enabling the development of secondary structures and larger bubble and spike amplitudes. Activating the SLR (RT₃) produces a similar qualitative improvement. However, although the AVSW calculation develops structure comparable to the AVSLR models, the stored kinetic energy in the vertical direction is clearly lower (Table 4). Among the SLR variants, RT₃, RT₅, and RT₆ are similar, with RT₃ and RT₆ exhibiting the highest kinetic energy, and thus preferred, closely followed by RT₅.

Model RT₄, which combines switches and SLR, reduces AV via both mechanisms. The resulting lower dissipation amplifies

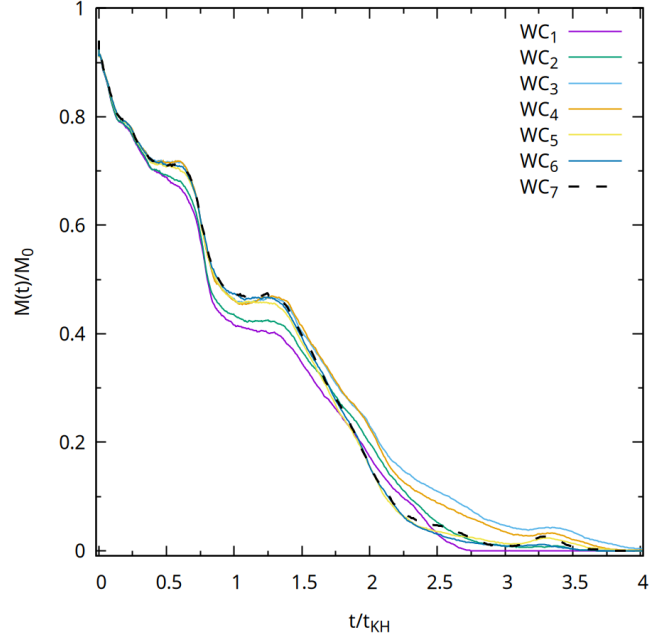


Fig. 8. Evolution of the cloud mass normalized to its initial value. The elapsed time is normalized to $t_{KH} = 2.9$.

numerical noise, leading to more asymmetric and more diffuse substructures. Despite the reduced viscosity, its vertical kinetic energy is the second-lowest in the set, reflecting the detrimental impact of noise on coherent RT growth.

3.6. The wind-cloud test

The destruction of a dense cloud embedded in a dilute supersonic wind is a demanding test that requires an accurate treatment of both shocks and shear-driven instabilities (Agertz et al. 2007, and references therein). That work showed that classic SPH underperformed relative to mesh methods, largely due to tensile instability that permeated SPH codes at the time. Modern SPH schemes, however, can handle this scenario well and yield competitive results (Frontiere et al. 2017; Wadsley et al. 2017; García-Senz et al. 2022). While the mitigation of E0-errors for this problem is more decisive than the choice of AV, the AV prescription still modulates the outcome and merits discussion.

The initial configuration follows Frontiere et al. (2017). The box size and particle count are given in Table 2. The wind and cloud have initial densities of $\rho = 1$ and 10, respectively, and are in pressure equilibrium at $P_0 = 1$. The cloud is initially at rest, while the wind moves supersonically with Mach number $\mathcal{M} = 2.7$. Wind and cloud particles are distributed in glass-like lattices, and the interface is sharp. Numerical studies have identified different hydrodynamic regimes as a function of the KHI characteristic time $t_{KH} = 2.9$ (corresponding to our setup⁶) at which the KHI dominates:

- $t \lesssim t_{KH}$. Direct impact and bow-shock formation control the initial stripping of the material from the sphere. According to Fig. 8, the cloud evolution is largely insensitive to the AV scheme.

⁶ We take the characteristic t_{KH} time as in Agertz et al. (2007), which is defined as $t_{KH} = 1.6 t_{crush}$, where t_{crush} is the cloud crushing time, right after the formation of the bow-shock; beyond this, KHI erodes the cloud surface.

Table 5. Cloud destruction rate, $\mathcal{R} = \dot{M}/M_0$, at $t = 1.8 t_{KH}$ and $t = 2.0 t_{KH}$ for the WC models.

	WC ₁	WC ₂	WC ₃	WC ₄	WC ₅	WC ₆
$\mathcal{R}_{1.8t_{KH}}$	-0.105	-0.112	-0.139	-0.144	-0.152	-0.162
$\mathcal{R}_{2.0t_{KH}}$	-0.121	-0.119	-0.127	-0.152	-0.174	-0.214

- $t_{KH} \lesssim t \lesssim 3t_{KH}$. Vorticity (primarily from KHI) peels and mixes cloud material into the wind. As seen in Fig 8, the removal occurs at a remarkably steady rate. The destruction rate $\mathcal{R} = \dot{M}(t)/M_0$ differs across AV schemes and provides a convenient ranking criterion (Table 5).
- $t \geq 3t_{KH}$. A complete dissolution of the cloud is challenging to model physically and is also resolution-dependent, since mass remnants approach zero at a sufficiently high resolution. In SPH, a small artificial thermal conductivity is commonly added to achieve complete mixing and thermodynamic equilibrium (Price 2008). By default, SPHYNX includes an artificial heat conduction term, following Price et al. (2018) with a diffusion coefficient $\alpha_u = 0.1$.

Overall, all AV schemes aim for complete cloud destruction at $t \leq 4t_{KH}$, in reasonable agreement with the results obtained with state-of-art mesh-codes such as ENZO (cited in Agertz et al. 2007). This reflects the strengths of our SPH flavor, namely, accurate gradient estimation via the integral approach (García-Senz et al. 2012) and the geometric density average force estimation (Wadsley et al. 2017; García-Senz et al. 2022). Still, the particular implementation of the AV is also important as it modulates the path to complete dissolution. The interval $1.25 t_{KH} \lesssim t \lesssim 2.5 t_{KH}$ is particularly relevant to our study because the slope of the $M(t)/M_0$ curves allows an easy classification of the AV variants.

According to Table 5, the highest destruction rates at $1.25 t_{KH} \leq t \leq 2.5 t_{KH}$ occur for WC₆, while the slowest are those with bare AV or with switches alone. Using this criterion, the AV schemes are ranked, from worst to best, in the same order as they appear in the table (we use the first row for ranking). However, note that SLR models without Balsara modulation (models WC₃ and WC₄) take longer to dissolve the cloud because they overestimate the effect of shocks, as seen in the Sedov test (Fig. 2); in the wind-cloud problem, this will delay efficient KHI-driven shredding.

4. Subsonic turbulence

Turbulence is ubiquitous across physical scales and is especially relevant in astrophysics, where it plays a major role in chemical mixing, stellar collapse, and supernova explosions, among other phenomena. High-resolution simulations of turbulence must be approached with care, as they are long and computationally expensive, with an adequate level of accuracy that must be maintained over many time steps. Subsonic turbulence is particularly demanding: even small accumulations of numerical error or a modest excess of dissipation can suppress the development of a faithful inertial range (i.e., the Kolmogorov cascade down to the dissipation scale). This has long challenged SPH approaches, which historically lagged behind Eulerian methods at low Mach numbers until recently.

Since the Reynolds number in SPH with standard AV scales roughly as $\text{Re} \propto \mathcal{M}$ at fixed resolution Price (2012), minimizing residual AV in smooth subsonic regions is essential to

sustain an inertial range. Modern SPH implementations that combine accurate gradient estimators, generalized volume elements, pairing-resistant kernels, and tighter dissipation control have now achieved Eulerian-quality results (Cabezón et al. 2025).

The simulation of subsonic turbulence presents an inverted dissipation problem relative to the previous tests in this work. In shocks, disks, and hydrodynamical instabilities, dissipation is driven primarily by strong compression and large velocity gradients. In contrast, subsonic turbulence is limited by the background numerical dissipation floor: if this baseline dissipation is even slightly too high, the energy cascade is prematurely damped, and the inertial range collapses. To mitigate spurious triggering by small-scale particle noise, it is therefore necessary to control the amplitude of the dissipation. To this end, we reuse the Balsara coefficients (already employed to modulate the SLR terms) to multiply the linear part of the AV, clamped from below by a floor, \mathcal{F} ,

$$\Pi_{ab} = \frac{-\max(\mathcal{F}, \bar{\mathcal{B}}_{ab})\alpha\bar{c}_{ab}\mu_{ab} + \beta\mu_{ab}^2}{\bar{\rho}_{ab}}, \text{ for } \mathbf{v}_{ab} \cdot \mathbf{r}_{ab} < 0, \quad (27)$$

with $0 \leq \mathcal{F} \leq 1$ and $\bar{\mathcal{B}}_{ab} = \frac{1}{2}(\mathcal{B}_a + \mathcal{B}_b)$. The clamp reflects our goal of a general-purpose dissipation scheme: a floor value of $\mathcal{F} = 1$ recovers the scheme used in previous tests. Allowing $\mathcal{F} \rightarrow 0$ yields excellent results in subsonic turbulence, but at the cost of effectively disabling AV in shear-dominated regions adjacent to shocks or in mixed regimes. We adopted the compromise of $\mathcal{F} = 0.5$: AV is sufficiently attenuated for subsonic turbulence, yet does not vanish in, for instance, shock-dominated flows. According to Table 7, this approach produces good results in all tests considered in this work.

Our metric to rank the turbulence models is to measure the location of the bottleneck peak (k_{max}) in the compensated velocity power spectrum. The works of Valdarnini (2016) and Cabezón et al. (2025) showed that SPH can reproduce both the inertial range and the bottleneck effect, provided that accurate gradients and sensitive AV control are used jointly. Therefore, given such a modern SPH implementation and a determined resolution, k_{max} becomes a proxy for effective dissipation: without amplitude control of the AV, dissipation acts earlier (lower k) and the bottleneck intrudes sooner, effectively reducing the inertial range. Hence, a larger k_{max} indicates a potential better model in this context.

Given the computational cost, we used the GPU-native SPH-EXA code to run at relatively high resolution, using $N = 800^3$ particles. The domain is a box of a side $L = 1$ with periodic boundary conditions and constant density $\rho = 1$. Particles start from a relaxed glass-like distribution. The driving is done via an Ornstein-Uhlenbeck process (Uhlenbeck & Ornstein 1930; Federrath et al. 2010), following a parabolic amplitude field $\sim 1 - (k - k_c)^2$, with $k = |\mathbf{k}|$, and $k_c = 2$, which corresponds to a length of $L/2$, and vanishing for $k \leq 1$ and $k \geq 3$. This approach is implemented to reduce the impact of the stirring on the energy cascade, facilitating a natural transfer of energy from larger to smaller scales. Our forcing uses a natural mixing of 1/3 compressive and 2/3 solenoidal contributions. All simulations begin in a homogenous state and are stirred until they become statistically stationary at $t \simeq 10$ with an RMS Mach number of $\mathcal{M} = 0.3$. The times are normalized to the sound-crossing time, which in our setup is $t_{sc} = L/c = 1$. We used a quasi-isothermal EOS with adiabatic index $\gamma = 1.001$. This setup is provided out-of-the-box in SPH-EXA and can be reproduced by modifying the AV as desired and using the following command line:




Fig. 9. Instantaneous compensated velocity power spectra at $t = 10$.

Table 6. Subsonic turbulence test: k values for the maximum of the power spectra of the velocity field at the bottleneck.

	T_1	T_2	T_3	T_4	T_5	T_6	T_7
k_{max}	8.68	25.40	46.88	59.04	32.99	41.67	49.47

`$./sphexa-cuda-init turbulence-prop turbulence -n 800 -s 10.0 -w 1.0,`

where `-n` sets N^{1D} , `-s` is the simulated physical time, and `-w` is the output frequency.

Figure 9 shows the resulting compensated power spectra at $t = 10$, and Table 6 lists the corresponding k_{max} values for models T_k in Table 2. As expected, AV with constant coefficients (T_1) activates too early, pushing the bottleneck effect to low k , near the box scale, leaving no inertial plateau and quickly leading to viscous roll-off. The best model combines AVSW with SLR (T_4) yielding nearly a decade-wide plateau and bottleneck peaking at $k = 59$. Pure SLR-based models, without AV amplitude modulation, achieve a very similar inertial range as (T_4), meaning that the large-to-intermediate dynamics are comparable. However, they show an earlier (and slightly stronger) encroachment of the bottleneck into the dissipation range. This is not surprising, since there is no mechanism to suppress AV in a smooth flow. Any residual compression or noise from particle disorder keeps a finite α active essentially all the time, leading to the earlier onset of dissipation. Model T_7 addresses this point, limiting the linear component of the AV, which delivers the second-best results.

5. Searching for the optimal AV scheme

Our goal in this section is twofold: to classify the different AV schemes from our test results, and to discuss specific drawbacks that may affect methods that do not directly control the value of the parameter α , such as SLR and Godunov-SPH.

5.1. Ranking the AV schemes

Given the spread of the results across the dissipation schemes in Table 2, the best method is the one that is most balanced over the entire test suite. Generally speaking, for shocks, the best

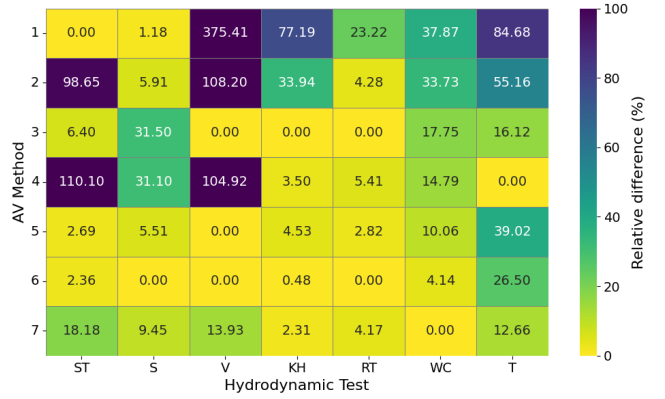


Fig. 10. Ranking heatmap of all AV methods for each test. Numbers show the relative difference to the best metric result overall for each test, using the values reported in Table 7. Lower values indicate better results.

schemes maintain a high AV at the shock front (e.g., models S_1 and ST_1 in Table 2). The opposite is true in regions subjected to strong shear, where AV should be kept low. On another note, the limitation in subsonic turbulence comes from the background random particle noise, which can spuriously trigger unnecessary AV. This needs to be controlled to allow for the emergence of the smaller-scale structure and the delay of the bottleneck and dissipation regimes.

Figure 10 shows a ranking heatmap of AV variants across all tests, including turbulence. The heatmap is indicative and consistent with the case-by-case discussion above and summarized in Table 7. The x -axis lists the tests, while the y -axis lists the AV schemes, indexed as in Table 7. Each cell is colored on a 0–100% scale representing the relative distance to the best metric result on that test. Quantitative metrics are those reported in the previous sections: L_1 values (ST, S, V), the perturbation growth rate (KHI), the vertical kinetic energy (RTI), the cloud-destruction rate in a fixed time interval (WC), and the bottleneck peak wavenumber in the compensated velocity power spectra (T). In each column (test), the best results have a value of 0%.

The horizontal accumulation of dark cells (many high percentages for a given method) in Fig. 10 indicates a poor combination of methodologies, while the light yellow bands indicate a robust scheme. For example, the constant AV scheme (number 1) is good for handling shocks owing to its larger dissipation capability, but it is not suitable for handling shear flows and turbulence. Using switches, AVSW (number 2), improves over scheme 1, but remains weak on average compared to the other methods. The AVSLR scheme (number 3) is the opposite, excels in simulating shear flows, but is worse with shocks. This is in agreement with the results of Sandnes et al. (2025), but also underscores the need for extra care in shock capturing and treatment of turbulence. Mixing switches and SLR in the AVSWSLR scheme (number 4) yields the best result in turbulence but degrades substantially in most other tests. The inclusion of the Balsara limiter to modulate the SLR, as suggested by Sandnes et al. (2025), AVSLRB (number 5), is able to cope with all fluid regimes successfully, except for turbulence. Our suggestion of using the square of the Balsara limiter in the SLR, AVSLRB2 (number 6) produces better results than method 5 and provides the best overall results in all tests, again with the exception of turbulence. Finally, adding the amplitude control of the AV, AVSLRB2 \mathcal{F} 0.5 provides a middle ground solution that behaves adequately overall, including subsonic turbulence.

Table 7. Quantitative information used in Fig. 10 to rank the first six tests cases with six AV schemes described in Table 2 and Sect. 3.

Test	ST		S	V	KHI	RTI	WC	T
	$L_1(v_x)$ ($\times 10^{-3}$)	ρ_{max}	$L_1(v_{rad})$ ($\times 10^{-2}$)	$L_1(v_\phi)$ ($\times 10^{-2}$)	$A(t/t_{KH} = 1.5)$ ($\times 10^{-2}$)	KE_Y ($\times 10^{-3}$)	$\mathcal{R}_{1.8t_{KH}}$	k_{max}
1-AV	2.97	4.00	2.57	5.80	2.87	6.81	-0.105	8.68
2-AVSW	5.90	4.02	2.69	2.54	8.31	8.49	-0.112	25.40
3-AVSLR	3.16	4.54	3.34	1.22	12.58	8.87	-0.139	46.88
4-AVSWSLR	6.24	4.56	3.33	2.50	12.14	8.39	-0.144	59.04
5-AVSLRB	3.05	4.02	2.68	1.22	12.01	8.62	-0.152	32.99
6-AVSLRB2	3.04	4.03	2.54	1.22	12.52	8.87	-0.162	41.67
7-AVSLRB2 \mathcal{F} 0.5	3.51	4.07	2.78	1.39	12.29	8.50	-0.169	49.47

Notes. The same information is also shown in the last row for a configuration giving good results in the turbulence study (see Sect. 4).

Table 8. Calculated models in the oscillating membrane test (Fig. 11).

	AV	AVSW	AVSLRB2	AVSWSLR	AVSLRB*2	AVSLRB*2 \mathcal{F} 05
α (min;max)	(0.2;0.2)	(0.05;1.0)	(1.0;1.0)	(0.05;1.0)	(1.0;1.0)	(0.5;1.0)
modulation	–	–	\mathcal{B}^2	–	\mathcal{B}^{*2}	\mathcal{B}^{*2}

Notes. Balsara numbers \mathcal{B} and \mathcal{B}^* are given by Eqs. (8) and (28).

5.2. Comments on possible disadvantages of SLR schemes

Despite the good performance of the AVSLR schemes shown above, two potential difficulties are worth noting. First, methods that do not explicitly control the AV coefficient α can struggle with linear-wave propagation. Second, in SLR-based formulations, the viscous form is not guaranteed to be strictly positive-definite pointwise, so care is needed to respect the second law of thermodynamics in discrete form.

Accurate propagation of nondissipative linear waves in SPH requires a negligible effective viscosity (Cullen & Dehnen 2010). Otherwise, the wave profile smears and its amplitude decays over time. Puri & Ramachandran (2014a,b) analyzed alternatives to AV switches (e.g., Godunov–SPH) and found that such schemes can generally handle 1D linear waves less robustly than switch-based AV. An AV+SLR implementation could face similar issues: without an explicit, time-dependent control of α , residual compression or noise might keep dissipation active. A straightforward remedy that involves reducing α aggressively will raise the noise level and might compromise other tests. A comprehensive study is beyond our scope, but this caveat deserves attention.

Astrophysical waves, however, are often surface or volume waves whose amplitude A_m naturally decays with distance from the source as $A_m \propto r^{-1/2}$ and $A_m \propto r^{-1}$, respectively. This geometric decay can outweigh that due to excessive AV damping, making 2D/3D wavefields more forgiving than the 1D case.

Here, we simulated low-amplitude waves generated by a harmonically oscillating cylindrical piston of a radius of $R_p = 0.035$ and length of $L = 1$, with a period of $T = 0.05$ and maximum radial velocity of $v_r^m = \pm 0.16$ (see for details García-Senz et al. 2012). The box and particle number match the Sedov setup: initial $\rho = P = 1$ give a sound speed of $c_s = 1.29$ for $\gamma = 5/3$. Although the amplitude of the wave is quite small, $\delta\rho/\rho \leq 6\%$, it is still considerably larger than that of typical sound waves and, therefore, a low amount of dissipation is expected.

A complication is that the Balsara limiter, \mathcal{B} (Eq. (8)), used to modulate SLR, was not originally designed for wave




Fig. 11. Velocity profile of a surface wave at $t = 0.35$ calculated with the AV schemes in Table 8. The L_1 values obtained are also shown when compared to the analytical result. The $1/\sqrt{r}$ dependence of maximum amplitude with distance is also shown (dashed line).

propagation. In this problem, the oscillation of $\nabla \cdot \mathbf{v}$ with the wave period induces the oscillation of \mathcal{B} , with the result that \mathcal{B} stays close to unity most of the time. Consequently, the SLRB and SLRB2 schemes behave similarly to the constant- α AV baseline (scheme 1 in Table 1). Therefore, to make the limiter sensitive to oscillatory flows it is necessary to redefine the Balsara coefficient, adopting a frequency-aware variant. A simple way to achieve this is to consider

$$\mathcal{B}_a^* = \frac{|\nabla \cdot \mathbf{v}|_a}{|\nabla \cdot \mathbf{v}|_a + |\nabla \times \mathbf{v}|_a + 10^{-4} \frac{c_a}{h_a} + \frac{2\pi}{T_a}}, \quad (28)$$

where T is the known wave period (the piston period here). For high or low frequencies, we have $\mathcal{B}_a^* \rightarrow 0/\mathcal{B}_a$, respectively.

Figure 11 compares the radial velocity $v_r(r)$ (with $r = \sqrt{x^2 + y^2}$) for the AV variants listed in Table 8 against the known

analytic profile for waves propagating in a membrane (e.g., [Elmore & Heald 1985](#)). We report the L_1 error in each panel. As expected, the lowest-viscosity cases (first, second, and fourth panels) show more noise but yield more symmetric waveforms. Switch-based AV exhibits larger scatter, though the mean follows the analytic profile well. The SLR model using the generalized limiter \mathcal{B}^* (Eq. (28)) shows a comparable performance, with an even smaller L_1 error, although with a gradual loss of symmetry over time. This suggests that unlike strictly 1D linear waves (no geometric decay), diverging or converging waves in 2D or 3D can be represented reasonably well with AVSLRB*2 or its floored variant AVSLRB*2F05, without any conflict with the rest of our test suite.

Finally, because SLR alters the reconstruction entering the dissipation term, we must ensure that the discrete viscous heating remains non-negative on average. As pointed out by [Price & Laibe \(2020\)](#), SLR schemes do not guarantee pointwise positive-definite AV by construction: the viscous work appears through products such as $(\mathbf{v}_{a,i} - \mathbf{v}_{a,j}) \cdot (\mathbf{v}'_{a,i} - \mathbf{v}'_{a,j})$, where \mathbf{v}'_a is the reconstructed velocity from Eq. (11), and these products are not necessarily positive. This can lead to occasional local violations of the second law of thermodynamics; for example, as unphysical heating in rarefaction waves. In some Riemann solvers, entropy fixes can be introduced to mitigate such issues ([Helluy et al. 2010](#)).

In practice, the likelihood of $(\mathbf{v}_{a,i} - \mathbf{v}_{a,j}) \cdot (\mathbf{v}'_{a,i} - \mathbf{v}'_{a,j}) < 0$ in shock-dominated systems is very small in our implementation. This is due to the Balsara modulation used in \mathbf{v}'_a (Eq. (19)), which makes sign reversal unlikely for $\mathcal{B} \geq 0.5$. Furthermore, in the shock-tube test (Sect. 3.1) we verified that no spurious entropy was produced in the rarefaction wave when using linear reconstruction. When $\mathcal{B} \leq 0.5$, the probability of sign reversals increases but still remains low. We monitored the fraction of negative pairwise products during the KHI run (model KH₆ in Table 2), where we typically have $\mathcal{B} \lesssim 0.5$, and we found

$$\frac{\text{negative products}}{\text{total number of products}} \leq 6\%, \quad (29)$$

with most inversions confined to small, noise-dominated velocities.

In contrast, smooth subsonic flows with $\mathcal{B} \simeq 0$ are more prone to local, nonphysical entropy variations. We therefore tracked both the evolution of the entropy of each particle and the total entropy during the propagation of the waves produced by the oscillating piston described above. These waves are nearly linear, so we expect the total entropy to increase over time but by a small, almost marginal amount. The change in entropy of particle a with respect to its initial value $S_{a,0}$ and normalized to the specific heat is

$$\left(\frac{S_a - S_{a,0}}{c_v} \right) = \ln \left(\frac{u_a}{u_{a,0}} \right) - (\gamma - 1) \ln \left(\frac{\rho_a}{\rho_{a,0}} \right). \quad (30)$$

Figure 12 shows the distribution of $(S - S_0)/c_v$ at $t = 0.35$ for models AVSW and AVSLRB*2. Far from the piston ($r > 0.1$), the entropy variation fluctuates almost symmetrically around zero in both cases, indicating a negligible change in entropy. This is consistent with the behavior of the total entropy shown in Fig. 13 (excluding the piston region and normalized to the number of particles). We see that the evolution of $\left\langle \frac{S - S_0}{c_v} \right\rangle$ is similar for both schemes and remains two orders of magnitude below the per-particle values shown in Fig. 12. Thus, summing these positive and negative local entropy variations yields a small, positive

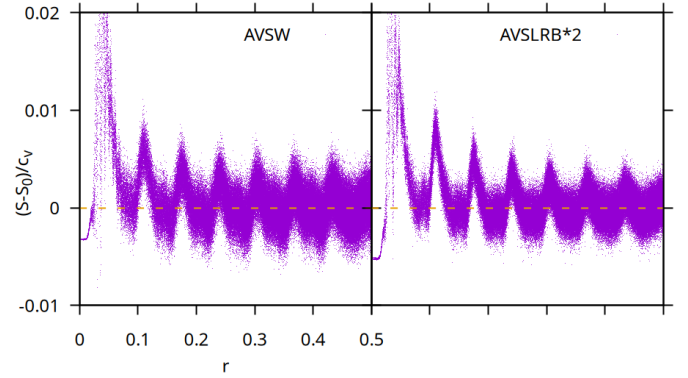


Fig. 12. Distribution of the entropy variation of each particle at $t = 0.35$ for models AVSW and AVSLRB*2 in Table 8.




Fig. 13. Evolution of averaged total entropy for models AVSW and AVSLRB*2 in Table 8.

residual (i.e. the expected slow increase of total entropy), with very similar trends in both calculations. We note that the AVSW curve shows a slight initial dip below zero, attributable to higher noise, before recovering.

6. Conclusions

[Price \(2024\)](#) and [Sandnes et al. \(2025\)](#) recently suggested that AV switches may be unnecessary if a slope-limited reconstruction (SLR) is used (see also [Frontiere et al. 2017](#); [Rosswog 2020a,b](#)). If confirmed, this assumption would generally represent a paradigm shift in the way shock-capturing and dissipation are implemented in SPH codes.

However, this does not come without a price. High, constant AV remains an excellent choice for shock capturing, although it does tend to overdamp turbulence, suppresses the emergence of hydrodynamic instabilities, and erase small-scale shear. Low, constant AV is great for turbulence, but produces strong post-shock oscillations. Switches may lag or trigger prematurely and leave residual noise in quiescent regions. Balsara limiters help in vortical flows, but they are unable to totally suppress unwanted dissipation in rotating shear flows. The bare SLR scheme reduces

false compression triggers, but does not provide adequate dissipation in strong shocks and fails to describe them. In other words, there is always a trade-off to consider when choosing a specific approach. Moreover, a bare SLR is not enough to completely control spurious AV triggers by random particle movements in smooth flows and it does not completely ensure the correctness of the sign in the particles' entropy change.

Even so, SLR is promising and, in our tests, it shows better results overall; in particular, for a Balsara-modulated SLR. In this work, we explore several SLR-based implementations across a broad test suite (shocks, shear flows, canonical instabilities, and subsonic turbulence) and suggest a balanced AV algorithm that performs robustly in all regimes. Our main findings are listed below:

- AVSLR schemes outperform our previous implementation of AVSW in SPHYNX across all tests, which is especially relevant for the study of subsonic turbulence;
- In line with Sandnes et al. (2025), a good multi-regime performance requires modulating the SLR reconstruction with, for example, the Balsara limiter \mathcal{B} , as shown in Eqs. (18) and (19). Although Sandnes et al. (2025) used the exponent $p = 1$, we find that $p = 2$ (our AVSLRB2) consistently yields better results, particularly in turbulence, without sacrificing shock handling;
- Nonphysical entropy behavior could appear if the bare SLR scheme is used (i.e., if SLR is not modulated). However, in regions subjected to mild or strong dissipation, the \mathcal{B} -modulation of SLR works to inhibit any nonphysical sign reversal in the entropy. Sign reversal can still occur in noise-dominated regions where switches also fail (e.g., Fig. 12);
- Combining switches with SLR (AVSWSLR) tends to degrade shock and Gresho–Chan performance, but it works very well for subsonic turbulence, as shown in Fig. 9 and Cabezón et al. (2025). This improvement is due to the additional reduction of the AV parameter α introduced by the switches;
- Motivated by the turbulence gains from reduced α , we tested a simple amplitude control mechanism of AV in our AVSLRB2 scheme, considering $\alpha \rightarrow \alpha \times \max[\mathcal{F}, \mathcal{B}]$, with $0 \leq \mathcal{F} \leq 1$ in Eq. (27). Taking $\mathcal{F} = 1$ recovers the AVSLRB2 scheme, while $\mathcal{F} = 0.5$ improves the turbulence spectra, while producing only minor changes in the other tests (see the last two rows of Table 7 and Fig. 10). This offers a switch-free path to near-optimal turbulence behavior without compromising shocks or instabilities.

In short, AVSLRB2 (SLR with Balsara modulation at $p = 2$) provides the best overall balance across shock-, shear-, and instability-dominated regimes. In addition, adding a modest AV amplitude floor ($\mathcal{F} = 0.5$) recovers much of the turbulence advantage seen with switches+SLR, without the broader downsides of switch-driven dissipation control. Therefore, we recommend using the AVSLRB2 \mathcal{F} scheme, setting the floor value to one in those cases where turbulence is not expected to be important, so that the scheme becomes AVSLRB2. In fluid regions where turbulence is relevant, it is more convenient to choose $\mathcal{F} = 0.5$.

An improvement of the technique would be to link the adopted value of \mathcal{F} to the local Reynolds number of the flow. Another future line of development is to explore the impact of other limiters in the SLR, different from the one used in this work (van Leer 1974), since Rosswoog (2024) recently showed that changing the limiter can modify the level of dissipation.

Returning to the question posed in the title as to whether switches still necessary, our results suggest that for the regimes

and resolutions explored here, time-dependent shock-trigger switches are likely not required to obtain robust dissipation control. In particular, SLR-based methods combined with continuous Balsara modulation (and, for subsonic turbulence, a modest amplitude floor) provide a single formulation that performs well from strong shocks to shear flows, instabilities, and low-Mach turbulence.

The results and conclusions presented here have been obtained using two state-of-the-art SPH codes, SPHYNX and SPH-EXA, together with a particular implementation of AV switches (Read & Hayfield 2012), which are also state-of-the-art schemes. Nonetheless, it would be valuable to repeat the same test suite with other modern SPH codes, such as SWIFT (Schaller et al. 2024) or PHANTOM (Price et al. 2018), which employ variants of the Cullen & Dehnen (2010) switch, to confirm (or further constrain) the advantages of SLR-based dissipation control over switch-based approaches.

Acknowledgements. We are very grateful to the referee, Daniel Price, both for constructive feedback and prompt reviewing of this work. This work has been supported by the Swiss Platform for Advanced Scientific Computing (PASC) project “SPH-EXA: Optimizing Smoothed Particle Hydrodynamics for Exascale Computing”. It has also been carried out as part of the SKACH consortium through funding from SERI. This work was also supported by a CHRONOS allocation from the Swiss National Supercomputing Centre (CSCS) under project ID ch18 on Alps. The authors acknowledge the support of the Center for Scientific Computing (sciCORE) at the University of Basel, where part of these calculations were performed.

References

- Agertz, O., Moore, B., Stadel, J., et al. 2007, *MNRAS*, **380**, 963
- Balsara, D. S. 1995, *J. Computat. Phys.*, **121**, 357
- Bauer, A., & Springel, V. 2012, *MNRAS*, **423**, 2558
- Cabezón, R. M., & García-Senz, D. 2024, *MNRAS*, **528**, 3782
- Cabezón, R. M., García-Senz, D., & Relaño, A. 2008, *J. Computat. Phys.*, **227**, 8523
- Cabezón, R. M., García-Senz, D., & Escartín, J. A. 2012, *A&A*, **545**, A112
- Cabezón, R. M., García-Senz, D., & Figueira, J. 2017, *A&A*, **606**, A78
- Cabezón, R. M., García-Senz, D., Seckin Simsek, O., et al. 2025, *A&A*, submitted, arXiv e-prints [arXiv:2503.10273]
- Cavelan, A., Cabezón, R. M., Grabarczyk, M., & Ciorba, F. M. 2020, in *PASC'20: Proceedings of the Platform for Advanced Scientific Computing Conference*, 11
- Cha, S.-H., & Whitworth, A. P. 2003, *MNRAS*, **340**, 73
- Cha, S.-H., Inutsuka, S.-I., & Nayakshin, S. 2010, *MNRAS*, **403**, 1165
- Cullen, L., & Dehnen, W. 2010, *MNRAS*, **408**, 669
- Elmore, W. C., & Heald, M. A. 1985, *Physics of Waves* (New York: Dover Publications, Inc.)
- Federath, C., Roman-Duval, J., Klessen, R. S., Schmidt, W., & Mac Low, M. M. 2010, *A&A*, **512**, A81
- Frontiere, N., Raskin, C. D., & Owen, J. M. 2017, *J. Computat. Phys.*, **332**, 160
- García-Senz, D., Cabezón, R. M., & Escartín, J. A. 2012, *A&A*, **538**, A9
- García-Senz, D., Cabezón, R. M., & Escartín, J. A. 2022, *A&A*, **659**, A175
- Gingold, R. A., & Monaghan, J. J. 1977, *MNRAS*, **181**, 375
- Godunov, S.-K. 1959, *Sb. Math.*, **47**, 271
- Gresho, P. M., & Chan, S. T. 1990, *Int. J. Numer. Methods Fluids*, **11**, 621
- Harten, A., Hyman, J. M., & Lax, P. D. 1976, *Commun. Pure Appl. Math.*, **29**, 297
- Helluy, P., Hérard, J.-M., Mathis, H., & Müller, S. 2010, *Comptes Rendus Mec.*, **338**, 493
- Inutsuka, S.-I. 2002, *J. Computat. Phys.*, **179**, 238
- Keller, S., Cavelan, A., Cabezón, R., Mayer, L., & Ciorba, F. 2023, in *Proceedings of the Platform for Advanced Scientific Computing Conference, PASC '23* (New York, NY, USA: Association for Computing Machinery)
- Lucy, L. B. 1977, *AJ*, **82**, 1013
- McNally, C. P., Lyra, W., & Passy, J.-C. 2012, *ApJS*, **201**, 18
- Monaghan, J. J. 1992, *ARA&A*, **30**, 543
- Monaghan, J. J. 1997, *J. Computat. Phys.*, **136**, 298
- Monaghan, J. J. 2005, *Rep. Prog. Phys.*, **68**, 1703
- Monaghan, J. J., & Gingold, R. A. 1983, *J. Computat. Phys.*, **52**, 374

- Morris, J. P., & Monaghan, J. J. 1997, *J. Computat. Phys.*, **136**, 41
- Osher, S., & Chakravarthy, S. 1984, *SIAM J. Numer. Anal.*, **21**, 955
- Price, D. J. 2008, *J. Computat. Phys.*, **227**, 10040
- Price, D. J. 2012, *MNRAS*, **420**, L33
- Price, D. J. 2024, arXiv e-prints [arXiv:2407.10176]
- Price, D. J., & Laibe, G. 2020, *MNRAS*, **495**, 3929
- Price, D. J., Wurster, J., Tricco, T. S., et al. 2018, *PASA*, **35**, e031
- Puri, K., & Ramachandran, P. 2014a, *J. Computat. Phys.*, **256**, 308
- Puri, K., & Ramachandran, P. 2014b, *J. Computat. Phys.*, **270**, 432
- Read, J. I., & Hayfield, T. 2012, *MNRAS*, **422**, 3037
- Rosswog, S. 2015, *MNRAS*, **448**, 3628
- Rosswog, S. 2020a, *ApJ*, **898**, 60
- Rosswog, S. 2020b, *MNRAS*, **498**, 4230
- Rosswog, S. 2024, *Comput. Model. Eng. Sci.*, **143** 1713
- Rosswog, S., & Diener, P. 2025, in *New Frontiers in GRMHD Simulations*, eds. C. Bambi, Y. Mizuno, S. Shashank, & F. Yuan, 235
- Sandnes, T. D., Eke, V. R., Kegerreis, J. A., et al. 2025, *J. Computat. Phys.*, **532**, 113907
- Schaller, M., Borrow, J., Draper, P. W., et al. 2024, *MNRAS*, **530**, 2378
- Sedov, L. I. 1959, *Similarity and Dimensional Methods in Mechanics*
- Sod, G. A. 1978, *J. Computat. Phys.*, **27**, 1
- Uhlenbeck, G. E., & Ornstein, L. S. 1930, *Phys. Rev.*, **36**, 823
- Valdarnini, R. 2016, *ApJ*, **831**, 103
- van Leer, B. 1974, *J. Computat. Phys.*, **14**, 361
- Wadsley, J. W., Keller, B. W., & Quinn, T. R. 2017, *MNRAS*, **471**, 2357
- Yee, H. C., Warming, R. F., & Harten, A. 1983, *Implicit Total Variation Diminishing (TVD) schemes for steady-state calculations*, NASA Technical Memorandum, NASA/STI Accession number: 19830014814