Mitigation of post-shock oscillations induced by artificial viscosity in discontinuous finite element methods

Guido Lodato, Luc Vervisch, Jean-Baptiste Chapelier

Computers and Fluids 241 (2022) 105491

Contents lists available at ScienceDirect

Computers and Fluids

journal homepage: www.elsevier.com/locate/compfluid

Mitigation of post-shock oscillations induced by artificial viscosity in discontinuous finite element methods

Guido Lodato [a,*], Luc Vervisch [a], Jean-Baptiste Chapelier [b]

[a] Normandie Université, INSA et Université de Rouen, 675 Avenue de l'Université, 76801 St. Etienne du Rouvray, France
[b] ONERA Aerospace Laboratory, 29 Avenue de la Division Leclerc, 92322 Chatillon, France

ARTICLE INFO

Keywords:
High-order methods
Spectral difference method
Shock-capturing methods
Shock detection
Post-shock oscillations

ABSTRACT

Current shock-capturing techniques for high-order discontinuous finite element methods based on modal energy shock sensors coupled to an artificial viscosity operator provide efficient and relatively simple solutions to tackle high-speed flows with discontinuities. Yet, due to inherent inhomogeneities in modal detection of shocks within elements, post-shock spurious oscillations can occur. These can be particularly severe in the case of slow-moving shocks. In the present paper, the origin of such oscillations is identified in the way the artificial viscosity is distributed across shocks by the modal shock detector. A novel approach is hence proposed to mitigate post-shock oscillations by introducing, via exponential averaging, a time delay in the artificial viscosity injection. Numerical one- and two-dimensional tests will show that a calibration of the time delay based on the shock propagation velocity can suppress, almost completely, post-shock oscillations, even in the most severe cases. Inspired by machine learning techniques, an efficient way to evaluate the propagation velocities of different shocks is also proposed, leading to a method to suppress post-shock oscillations, induced by the artificial viscosity, which is extremely simple, and relatively inexpensive.

1. Introduction

The present paper focuses on high-order discontinuous finite element (DFE) methods applied to high-speed flows in the presence of shocks. It is well known that discontinuous solutions can generate spurious numerical oscillations (Gibbs phenomena) which need to be mitigated to ensure numerical stability. Among the many available techniques to achieve this, artificial viscosities [1–5] are relatively popular for DFE thanks to their simplicity and ease of implementation, especially in the case of three-dimensional unstructured meshes.

To identify regions in the flow where the solution needs to be smoothed out by the introduction of artificial viscosity (AV), a sufficiently localized shock detector is adopted. Thanks to the spectral nature of DFE, shock detection can be easily achieved by evaluating the modal energy decay of a selected flow quantity in each element of the discretization. Relevant detection techniques include the evaluation of the ratio between the energy of the highest mode and the energy of the whole spectrum, or the computation of a least-square power fit of the modal energy spectrum [2–4]. However, supposedly due to inherent inhomogeneities in the way the modal sensor detects a discontinuity within the element [4], post-shock oscillations can occur, which can become particularly severe downstream of slow-moving shocks (see [6] and references therein). Since these oscillations are triggered by spurious perturbations related to the way shocks are detected, an increase in the amount of the injected artificial viscosity would make them stronger. On the other hand, if a simple augmentation of artificial dissipation is not a viable option, acting on the sensor itself can give some improvements and the problem can be somewhat mitigated by a suitable choice of the signals used to detect the discontinuities [7,8]. Yet, depending on the flow parameters, the improvements can remain relatively limited.

Restricting the attention to the spectral difference (SD) scheme [9–11], the main subject of the present study concerns an improvement to the AV strategy such as to mitigate post-shock oscillations. The proposed approach is readily applicable on a wide range of similar schemes, such as the Discontinuous Galerkin (DG) [12–14], the spectral volume (SV) [15,16] and the flux reconstruction (FR) [17,18] type schemes.

The paper is organized as follows: the adopted physical model and the relevant numerical solution method are briefly presented in Section 2; Section 3 presents the post-shock oscillation phenomenon with some discussion concerning the main causes; the proposed method to mitigate or suppress post-shock oscillations is detailed in Section 4 and one of the relevant main ingredients, namely, the evaluation of the propagation velocities of each shock, is formalized in Section 5;

* Corresponding author.
E-mail address: guido.lodato@insa-rouen.fr (G. Lodato).

https://doi.org/10.1016/j.compfluid.2022.105491
Received 29 June 2021; Received in revised form 30 March 2022; Accepted 27 April 2022
Available online 4 May 2022

finally, the relevant results on one- and two-dimensional test cases are discussed in Section 6 and some concluding remarks are made in Section 7.

2. Governing equations and numerical setup

The SD solver utilized in the present study is based on the three-dimensional Navier–Stokes equations for an ideal, Newtonian fluid. Although all the relevant test cases are one- and two-dimensional, the full three-dimensional formulation is presented below, the reduction in dimensionality being obtained, in a straightforward way, as a particular case of the presented equations.

Let ρ, $\mathbf{u}$ and E be the fluid's density, the velocity vector and the total energy (internal + kinetic), respectively. The three-dimensional Navier–Stokes equations in conservative form are:

$$\frac{\partial \mathbf{U}}{\partial t} + \nabla \cdot \mathbf{F}_c + \nabla \cdot \mathbf{F}_v = \mathbf{0}, \tag{1}$$

where $\mathbf{U} = \begin{pmatrix} \rho & \rho\mathbf{u} & \rho E \end{pmatrix}^T$ is the vector of conservative variables, t is the time, and the convective and viscous fluxes are, respectively,

$$\mathbf{F}_c = \begin{pmatrix} \rho\mathbf{u} \\ \rho\mathbf{u} \otimes \mathbf{u} + p\mathbf{I} \\ (\rho E + p)\mathbf{u} \end{pmatrix}, \quad \text{and} \quad \mathbf{F}_v = \begin{pmatrix} \mathbf{0} \\ -2\mu\mathbf{D} - \lambda\,\mathrm{tr}(\mathbf{D})\mathbf{I} \\ -2\mu\mathbf{D}\cdot\mathbf{u} - \lambda\,\mathrm{tr}(\mathbf{D})\mathbf{u} - k\nabla T \end{pmatrix}, \tag{2}$$

the symbol $\otimes$ representing the tensor product operator. In the above equations, p is the pressure, $\mathbf{I} \in \mathbb{R}^{3\times3}$ is the identity matrix, μ is the dynamic viscosity, λ is the second coefficient of viscosity which is set according to the Stokes' hypothesis, k is the thermal conductivity, T is the temperature, and $\mathbf{D}$ is the rate of deformation tensor, namely,

$$\mathbf{D} = \tfrac{1}{2}\left(\nabla\mathbf{u} + \nabla\mathbf{u}^T\right), \quad \text{with} \quad \mathrm{tr}(\mathbf{D}) = \nabla\cdot\mathbf{u}. \tag{3}$$

Assuming the validity of the ideal gas law, the total energy E is related to the pressure and temperature as

$$\frac{p}{\rho} = (\gamma-1)c_v T, \quad \rho E = \frac{p}{\gamma-1} + \frac{1}{2}\rho\mathbf{u}\cdot\mathbf{u}, \tag{4}$$

where $\gamma = c_p/c_v$ is the ratio between specific heat capacities at constant pressure and constant volume. In the case of Euler (or inviscid) computations, the viscous flux $\mathbf{F}_v$ is set to zero everywhere in the flow, with the exception of the regions where shocks and discontinuities are detected, in which case, only the artificial viscosity (AV) from the shock capturing procedure remains active.

The above equations are solved using the high-order SD method for unstructured spatial discretization [9–11,16,19,20]. Details about the relevant implementation on unstructured hexahedral grids are summarized in Ref. [7, Appendix A]. Unless stated otherwise, all the results reported in the present study have been computed using the Roe flux with entropy fix [21,22] for the inviscid interface fluxes, whereas the interior penalty (IP) flux is adopted for the viscous fluxes. Time integration is performed with the RK45-SSP Runge–Kutta scheme [23] with an adaptive time step set at 40% of the admissible time step to satisfy the Courant–Friedrichs–Lewy (CFL) and Fourier conditions for the convective and diffusive fluxes, respectively.

In the case of high-speed flows with shocks and discontinuities, the above equations are augmented by the introduction of an AV term. The selected method, originally developed for the DG scheme [2,3], combines a highly selective spectral sensor, based on the modal decomposition via orthogonal polynomials, with a consistently discretized AV. The original method involved the use of either a *Laplacian* or a *physical* viscous term [2]. When the former is adopted, Eq. (1) is augmented with a Laplacian diffusion term $\nabla\cdot(\varepsilon\nabla\mathbf{U})$, where ε represents the added AV to be introduced in regions where the flow is under-resolved, such as across shocks and contact discontinuities. For the physical AV term, an entropy preserving formulation is adopted [5]. In this case, an additional viscous term $\nabla\cdot\mathbf{F}_\varepsilon$ is introduced in Eq. (1) with

$$\mathbf{F}_\varepsilon = \begin{pmatrix} \mathbf{0} \\ -\rho\varepsilon\,\mathrm{tr}(\mathbf{D})\mathbf{I} \\ -k_\varepsilon\nabla T \end{pmatrix}, \tag{5}$$

where the artificial thermal conductivity is computed as $k_\varepsilon = \rho\varepsilon c_p/\mathrm{Pr}_\varepsilon$ and Pr_ε is set constant and equal to the gas Prandtl number.

Concerning the modal sensor, a characteristic-based shock detection formulation is adopted [7,8]. Accordingly, the AV is injected based on the worst resolved signal among the density and the acoustic ones. The acoustic signals, in particular, are evaluated by projecting the solution on the left eigenvectors of the conservative flux Jacobian computed using the average solution $\overline{\mathbf{U}}$ within the element [7, Appendix D]. The element-wise AV is computed, one direction at a time, based on the modes of the selected signals along the three orthogonal directions defining the standard cubic element in computational space. The maximum value of the AV is then retained.

In the present implementation, as detailed in Ref. [7], the modal sensor is computed, inside each element, as

$$s_e(\psi) = \log_{10}\left[\frac{\hat{\psi}_n^2}{\sum_{i=1}^n \hat{\psi}_i^2}\right], \tag{6}$$

where ψ is the selected signal (density and/or acoustic), $\hat{\psi}$ the relevant modes of the Legendre polynomial basis and n the discretization order. The value of the artificial viscosity ε_e is then evaluated as

$$\varepsilon_e(s_e) = \begin{cases} 0 & \text{for } s_e < s_0 - \kappa, \\ \frac{\varepsilon_0}{2}\left[1 + \sin\frac{\pi(s_e - s_0)}{2\kappa}\right] & \text{for } s_0 - \kappa \le s_e \le s_0 + \kappa, \\ \varepsilon_0 & \text{for } s_e > s_0 + \kappa, \end{cases} \tag{7}$$

where ε_0, s_0 and κ are additional parameters to be chosen empirically. In particular, the AV intensity scaling factor ε_0 is computed from a scaling velocity and the element size h as

$$\varepsilon_0 = C_\varepsilon \lambda_{\max} h/(n-1), \tag{8}$$

where C_ε is a user-defined coefficient, commonly set equal to one, and $\lambda_{\max}$ is the maximum wave speed (*i.e.*, the spectral radius of the inviscid flux Jacobian) in the whole domain [24]. Concerning the threshold s_0 and sensor tolerance κ, these are computed via a self-calibration algorithm through a manufactured solution [7,25].

Finally, to guarantee numerical stability in the case of strong shocks, the shock-capturing method is coupled with a high-order, conservative, positivity-preserving scheme [26]. The generalization to the SD scheme is rather straightforward and the relevant details regarding the actual implementation are reported in Ref. [7, Appendix C].

3. The post-shock oscillations problem

To highlight the problem of post-shock oscillations (PSO) with the present numerical setup, a one-dimensional test involving a shock traveling across the domain at a fixed velocity D is used. It shall be noticed that relatively mild PSO are also present, for instance, in the classical Sod tube Riemann problem [4,7,27]. The present test, on the other hand, is designed to exacerbate the issue and hence the computation is performed in a moving relative reference frame, such that the propagation velocity of the shock in the reference frame of the mesh can be set at a given (small) percentage of the expected shock speed in the fluid at rest.

For a stationary ($D = 0$) shock, Rankine–Hugoniot conditions relate the left and right states (subscript L and R, respectively) as

$$\frac{\rho_R}{\rho_L} = \frac{(\gamma+1)\mathrm{M}^2}{(\gamma-1)\mathrm{M}^2+2}, \quad \frac{u_R^0}{u_L^0} = \frac{\rho_L}{\rho_R}, \quad \frac{p_R}{p_L} = \frac{2\gamma\mathrm{M}^2 - \gamma + 1}{\gamma+1}, \tag{9}$$

with M the shock Mach number and $u_L^0 = \mathrm{M}\sqrt{\gamma p_L/\rho_L}$. The 0 superscript has been used to denote velocities for the steady shock case. On the other hand, the propagation velocity of the shock in the fluid at rest, denoted as D^*, is obtained in a reference frame moving with velocity $U_{\mathrm{ref}} = u_L^0$ (left state at rest), namely, $D^* = -u_L^0$. In order to obtain a shock which propagates at a speed σD^*, the reference frame shall move with a translation velocity equal to $U_{\mathrm{ref}} = -\sigma D^* = \sigma u_L^0$. Considering,

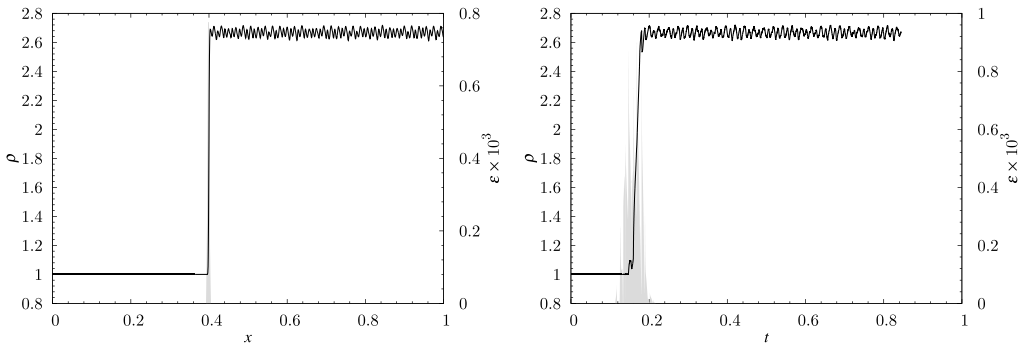

Fig. 1. Slowly moving Mach 2 shock: solid line, density (left axis); gray filled line, shock-capturing viscosity (right axis).

Table 1
PSO amplitude ($\%\Delta\rho$) of moving shocks. Sensor: ρ|ac, density-acoustic sensor [7]; ρ, density sensor.

Mach	Sensor	PSO amplitude ($\%\Delta\rho$)				
		$5\%D^*$	$30\%D^*$	$50\%D^*$	$70\%D^*$	$90\%D^*$
1.1	ρ|ac	0.46	0.35	0.12	0.08	0.06
2.0	ρ|ac	5.94	1.06	0.39	0.18	0.06
	ρ	5.70	1.50	0.67	0.30	0.15
3.5	ρ|ac	5.55	1.07	0.17	0.06	0.01
5.0	ρ|ac	6.47	1.29	0.18	0.10	0.00
9.0	ρ|ac	6.14	1.12	0.09	0.01	0.00

for instance, a M = 2 shock, setting $\rho_L = 1$, $p_L = 1$, the initial states, in the relative mesh reference frame, are readily obtained:

$$u_L = u_L^0 - U_{\text{ref}} = (1-\sigma)2.366,$$
$$u_R = u_R^0 - U_{\text{ref}} = 0.887 - \sigma 2.366, \quad \rho_R = 2.667, \quad p_R = 4.5.$$

Fig. 1 shows the density profile obtained on a unitary domain ($L = 1$), with 400 5th-order elements, after the shock has traveled a fixed distance $0.2L$ in the case that $\sigma = 5\%$, whereas Table 1 summarizes the amplitudes of PSO in density,[1] as a percentage of the density jump, for different shock velocities and different Mach numbers. Notice that, unless stated otherwise, for all the computations in this and the next section, the Laplacian model with the density-acoustic sensor is adopted, scaling the AV to the spectral radius of the inviscid flux Jacobian with a unitary scaling factor C_ϵ (*cf.* Eq. (8)).

Focusing on the M = 2 shock with $\sigma = 5\%$, the density profile downstream the shock presents significant oscillations in density whose amplitude is about 6% of the density jump. For this extremely high level of PSO, the type of sensor used has marginal effects on the PSO amplitude (*cf.* Table 1). Notice that the AV is quite localized around the shock with a maximum value of about 0.8×10^{-3}. Yet, as it can be observed in the right plot, which show the time history at a fixed location, the evolution of ϵ by the time the shock passes by is quite intermittent.

As pointed out in Ref. [4], PSO are (supposedly) due to inherent inhomogeneities in the modal detector or, in other words, due to differences in the way the same discontinuity is detected depending upon its position within the elements. The effect of such inhomogeneities can be mitigated by a more suitable distribution of the AV across the shock [7] and the shock-sensor plays a key role in the amount of PSO. In order to verify this statement, a very simple test is to eliminate any inhomogeneity in ϵ. Indeed, by just setting $\epsilon = 0.8 \times 10^{-3}$ everywhere in the domain the density profile, shown in Fig. 2, is completely free from PSO.

[1] The PSO amplitude is evaluated as 4σ, where σ is the standard deviation of the time history of the density at a fixed location, after the shock has passed.

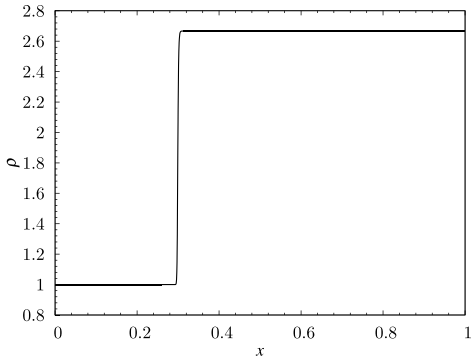

Fig. 2. Slowly moving Mach 2 shock with uniform and constant ϵ over the computational domain.

Although a uniform distribution of the AV in the whole domain is, of course, rather pointless, the present test suggests that the inhomogeneities and the local distribution of the AV, while the shock travels through the elements, play a central role on the onset of PSO. In the next section, an approach to improve the behavior of the AV will be presented.

4. The exponential averaging approach

As pointed out in the previous section, two main factors appear to be the culprit of PSO behind moving shocks: (a) the way the shock-sensor is activated in the neighborhood of the shock, which determines the AV distribution across the shock and (b) the inhomogeneities in shock detection when the shock travels across the element, which may introduce a spurious forcing in the flow.

The former aspect pertains the spatial distribution of the AV and whatever solution might be envisaged as a mean of improvement is most likely to have a significant impact on the locality of the shock capturing approach (*e.g.*, some sort of smoothing across the neighbor cells).

The latter aspect, on the other hand, can be addressed with a completely local approach, inside each element, by constraining the AV to *respond less promptly* to changes in the shock sensor while the shock travels across the element. In other words, a suitably calibrated *time delay* can be introduced such as to prevent the element-wise value of AV from changing too rapidly within the time the shock propagates through the element.

To this end, a time memory of the past values of the AV in every element is introduced by computing the AV value via an *exponential moving average* procedure (or *exponential smoothing*). Let ϵ_e^i be the computed value of AV in the eth element at the ith iteration. Then,

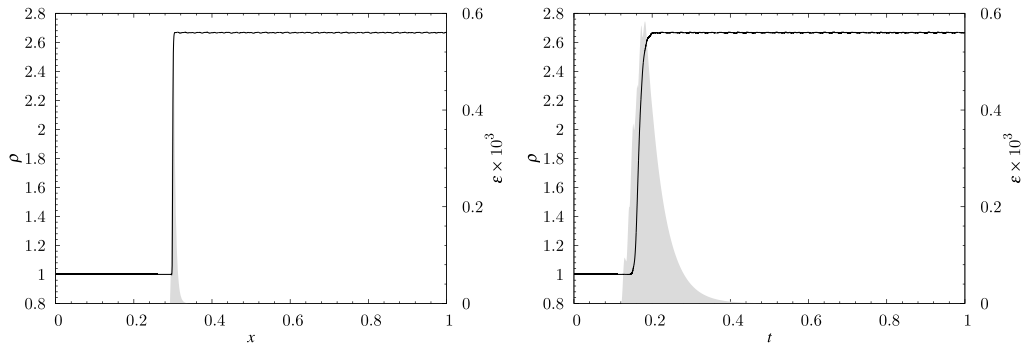

(a) spatial distribution (b) probe at $x = 0.48$

Fig. 3. Slowly moving Mach 2 shock with exponential smoothing at imposed shock speed (*cf.* Eq. (15)): solid line, density (left axis); gray filled line, shock-capturing viscosity (right axis).

Table 2
PSO amplitude with exponential smoothing (Mach 2.0, $C_{ea} = 0.5$).

D/D^*	0.05	0.30	0.50	0.70	0.90
PSO amplitude (%$\Delta\rho$)	0.397	0.059	0.007	0.001	0.000
PSO reduction (%)	93.3	94.4	98.2	99.6	100.0

the exponential averaging algorithm to update the value of ε_e^i at every iteration and in every element is:

$$\langle\varepsilon\rangle_e^i = \begin{cases} \varepsilon_e^0, & \text{for } i = 0, \\ \alpha\varepsilon_e^i + (1-\alpha)\langle\varepsilon\rangle_e^{i-1}, & \text{for } i \geq 1, \end{cases} \tag{10}$$

where the *smoothing factor* $\alpha \in [0 : 1]$ determines how much memory the AV has of its previous values in the element and the angled brackets have been introduced to distinguish between the instantaneous value, ε_e^i, and the exponentially averaged one, $\langle\varepsilon\rangle_e^i$. Clearly, the higher the value of α, the shorter the memory of the past: a higher α discounts older values faster.

The time constant τ of the exponential moving average (*i.e.*, the time delay of the averaged response of a unit step function to reach 63.2% of the original signal) can be related to the smoothing factor via the relation

$$\alpha = 1 - \exp(-\Delta\theta/\tau), \tag{11}$$

where $\Delta\theta$ is the sampling time interval between subsequent values of $\langle\varepsilon\rangle_e^i$, which is set equal to the time step divided by the number of stages, n_{RK}, of the Runge–Kutta scheme ($\langle\varepsilon\rangle_e^i$ is re-evaluated at each Runge–Kutta stage), namely,

$$\Delta\theta = \Delta t/n_{RK}. \tag{12}$$

If D is the propagation velocity of the shock and Δx is a measure of the cell size, then the time for the shock to traverse the element is estimated as

$$\tau_D = \Delta x/D, \tag{13}$$

which establishes a lower bound for the time constant τ of the exponential moving average. In other words, changes in $\langle\varepsilon\rangle_e^i$ shall happen on a time scale which is of the order of (or longer than) the time needed for the shock to cross the whole element:

$$\tau \geq \tau_D, \quad \text{or} \quad \tau = \tau_D/C_{ea}, \tag{14}$$

where C_{ea} is a dimensionless coefficient of order unity. From Eqs. (11)–(14), the smoothing factor becomes:

$$\alpha = 1 - \exp\left(-C_{ea}\frac{D\Delta t}{n_{RK}\Delta x}\right), \quad \text{with} \quad C_{ea} \leq 1. \tag{15}$$

The M = 2 test performed in the previous section, with exponential smoothing and the smoothing factor computed from Eq. (15) using the exact value $D = 5\%D^* = 0.118$ and $C_{ea} = 0.5$ (more on the optimal value of C_{ea} later), produces the results depicted in Fig. 3. As it can be seen, the AF distribution is slightly wider around the shock and, most importantly, its evolution in time is much smoother, leading to a PSO amplitude of about 0.4% of the density jump. As a side effect of the introduced delay in the AV, its maximum value is slightly lower overall and—although difficult to see in the plots—the shock is a bit sharper (yet completely free from Gibbs phenomena).

PSO amplitudes and reductions (with respect to similar computations without exponential smoothing) for different values of the shock propagation velocity are reported in Table 2. The relevant reductions range from a minimum of about 93%, for slow-moving shocks, to almost 100%, for fast-moving shocks. In practice, for the present Mach number and discretization, PSO are completely suppressed for shock traveling beyond 70% of their propagation speed in fluid at rest D^*. Notice, however, that, at such relatively high displacement speeds, PSO are extremely mild (less than 0.2%$\Delta\rho$) even without exponential smoothing (*cf.* Table 1).

4.1. The C_{ea} parameter sensitivity

The C_{ea} parameter sets the time constant of the exponential smoothing with respect to the shock displacement time within the element. In particular, from Eq. (14), the lower the value of C_{ea}, the slower the time response of the exponentially averaged AV in each cell. Notice that a slow response also means a smoother AV distribution and, indeed, small values of C_{ea} promote stronger reductions in PSO. At the same time, a slow response also means that (a) the AV might not be able to reach sufficiently high values and, at the same time, (b) the AV is less localized in the neighborhood of the shock because of the longer delay to drop back to zero after the shock has passed.

To better clarify the behavior of $\langle\varepsilon\rangle_e^i$, its time history recorded at a fixed point in space is plotted in Fig. 4 for different values of C_{ea}, ranging from 0.1 to its upper unitary bound. Notice that, to have a perception of the spatial distribution of the AV across elements, the transformed coordinate $x = Dt$ is represented on the top axis, with D the shock displacement speed and the tick marks spaced according to the elements' width.

As anticipated, the C_{ea} parameter has an impact on both the peak value and the width of the AV signal. Concerning the former, occasional very high values of AV, which can be observed without exponential smoothing (*cf.* Fig. 4b), are completely avoided regardless of the strength of exponential smoothing. The peak of AV is marginally reduced for $C_{ea} \in [0.5 : 1.0]$ (about −6% at most). On the other hand, for values of C_{ea} smaller than 0.5, significant reductions in the AV peak are observed (*e.g.*, −35% for $C_{ea} = 0.1$). With regards to

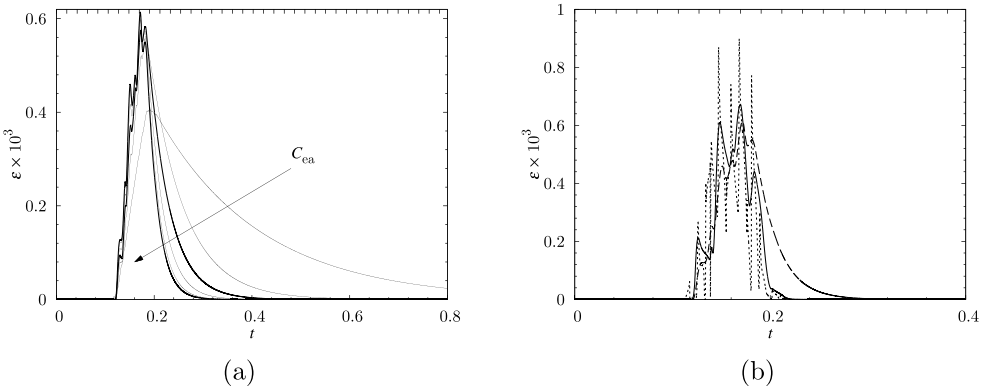

Fig. 4. Time history of $\langle\varepsilon\rangle_e^i$, at a fixed location for the slowly moving Mach 2 shock with exponential smoothing at imposed shock speed (*cf.* Eq. (15)). Plot (a): thin solid line, $C_{ea} \in [0.1 : 0.9]$ with 0.2 step increments; thick solid line, $C_{ea} = 0.5$ and 1.0. Plot (b): solid line, $C_{ea} = 3.0$; dashed line, $C_{ea} = 1.0$; dotted line, no exponential smoothing (the top axis shows the equivalent element width using the transformation $\Delta x = D\Delta t$).

Table 3
PSO amplitude (PSO.a) and reduction (PSO.r) for the slow-moving shock with $D = 5\%D^*$: L, Laplacian AV; P, physical AV.

$n \times N_e$	M	AV	PSO.a ($\%\Delta\rho$)/PSO.r (%)			
			$C_{ea} = 0.1$	$C_{ea} = 0.5$	$C_{ea} = 1.0$	$C_{ea} = 3.0$
5 × 200	1.1	L	0.178/61.5	0.233/49.7	0.219/52.8	
	2.0	L	0.608/89.8	1.965/67.0	2.034/65.8	
	2.0	P	1.331/74.0	1.222/76.1	1.438/71.9	
	5.0	L	0.467/90.7	0.178/96.5	1.018/79.9	
5 × 400	1.1	L	0.050/89.2	0.115/75.1	0.170/63.1	0.207/55.0
	2.0	L	0.046/99.2	0.393/93.3	0.907/84.6	3.293/44.2
	2.0	P	0.200/96.0	0.187/96.3	0.443/91.2	1.397/72.2
	5.0	L	0.402/92.3	0.178/96.6	1.072/79.4	4.212/19.1
7 × 400	1.1	L	0.050/90.9	0.093/82.8	0.120/77.9	
	2.0	L	0.070/98.4	0.556/87.2	0.922/78.8	
	2.0	P	0.204/94.8	0.116/97.0	0.244/93.8	
	5.0	L	0.047/98.7	0.089/97.5	0.233/93.5	
5 × 600	1.1	L	0.049/89.4	0.116/75.0	0.172/62.8	
	2.0	L	0.030/99.5	0.396/93.3	1.042/82.4	
	2.0	P	0.048/99.0	0.174/96.5	0.437/91.3	
	5.0	L	0.373/92.8	0.179/96.6	1.081/79.1	

the spatial distribution, considering as a baseline test the simulation without exponential smoothing, where the AV signal extends across about 5 elements, the exponential smoothing widens the AV signal, downstream the shock only, of about 60%, for $C_{ea} = 1.0$, and 160%, for $C_{ea} = 0.5$. Significantly wider AV distributions are obtained for $C_{ea} < 0.5$.

In terms of PSO (shown in the next section), the lower the value of C_{ea}, the higher their reduction. For this Mach number, PSO in density are between 0.4% and 0.9% of the density jump for $C_{ea} \in [0.5 : 1.0]$ (*i.e.*, a PSO reduction between 85% and 93%) and down to 0.05% of the density jump for $C_{ea} = 0.1$ (*i.e.*, a PSO reduction of about 99%).

Finally, for the relatively large value of $C_{ea} = 3.0$ (*cf.* solid line curve in Fig. 4b), marginal smoothing is achieved on $\langle\varepsilon\rangle_e^i$, with practically no change in AV localization. Concerning PSO, at this value of the smoothing parameter, their amplitude remains as high as 3.3% of the density jump, that is to say, a reduction of *only* about 44% compared to the baseline test without exponential smoothing (*cf.* Table 3).

4.2. Sensitivity to mach, order and grid resolution

The slow moving shock test ($D = 5\%D^*$) has been performed for different values of the shock Mach number (M = 1.1, 2.0 and 5.0), different discretization orders ($n = 5$ and 7) and different numbers of elements ($N_e = 200$, 400, 600). For all the computations, the shock sensor is the density-acoustic one and the exponential smoothing parameter C_{ea} takes the values 0.1, 0.5 and 1.0. The higher value $C_{ea} = 3.0$ is also tested for $n = 5$ and $N_e = 400$. As pointed out in the previous section, this value produces mild smoothing without any appreciable penalty in AV localization. For every test case, PSO reductions in density are measured with respect to the baseline computation performed without exponential smoothing. The relevant results are summarized in Table 3.

Fig. 5 shows the results in terms of PSO amplitude, as a percentage of the density jump, and PSO percentage reduction with respect to the baseline computations without exponential smoothing, when the Laplacian AV model is adopted. As anticipated, overall, the lower is the value of C_{ea}, the higher is the reduction in PSO. There are some exception though, as, for example, the fifth-order computations at Mach 5.0, for which the case with $C_{ea} = 0.5$ performs better than the case with $C_{ea} = 0.1$ ($\sim$ 97% versus $\sim$ 92% reduction, respectively). For values of $C_{ea} \in [0.5 : 1.0]$, thus without an excessive penalization in terms of AV localization (*cf.* Section 4.1), PSO reductions are quite significant. With the exclusion of the lowest Mach number tests on the coarse mesh ($N_e = 200$), for which the PSO amplitudes are anyway extremely small and of the order 0.2% of the density jump, PSO reductions are more than 60%. With the sole exception of the test at M = 2 on the coarse mesh (about 2% amplitude and 67% reduction), PSO amplitudes are less than 0.6% of the density jump for $C_{ea} = 0.5$, which appears to give a good balance between PSO reduction and AV localization.

Switching to the entropy preserving physical AV model [5], a comparison with the Laplacian model is shown in Fig. 6 for the Mach 2.0 case. As it can be noticed, for $C_{ea} \geq 0.5$, the physical viscosity model has a clear advantage over the Laplacian one, both in terms of PSO reduction and PSO amplitude. This is possibly due to the physical formulation being less susceptible of developing additional unphysical perturbations, such as overheating errors and AV-induced inconsistencies between mass, momentum and energy transport equations. The results degrade slightly when $C_{ea} < 0.5$, for which the excessive delay in the AV and the consequent reduction in the peak value, promotes the onset of Gibbs phenomena in the shock profile (not shown).[2]

Finally, when C_{ea} is set equal to 3.0 to favor AV localization, PSO reductions are much less pronounced, especially at the higher Mach numbers. At M = 5.0, for instance, PSO reductions go from about 97%, for $C_{ea} = 0.5$, down to about 19%, when $C_{ea} = 3.0$.

[2] It is worth noticing that the Laplacian and the physical models have been compared with the same AV intensity scaling factor. In other words, tests have been performed scaling the AV to the spectral radius of the inviscid flux Jacobian with a unitary coefficient for both models (*cf.* Eq. (8)). Under these circumstances, the increased dissipation brought by the Laplacian formulation, makes this last more robust to further reductions in the computed AV (but also produces shock profiles which are less sharp than those obtained with the physical AV).

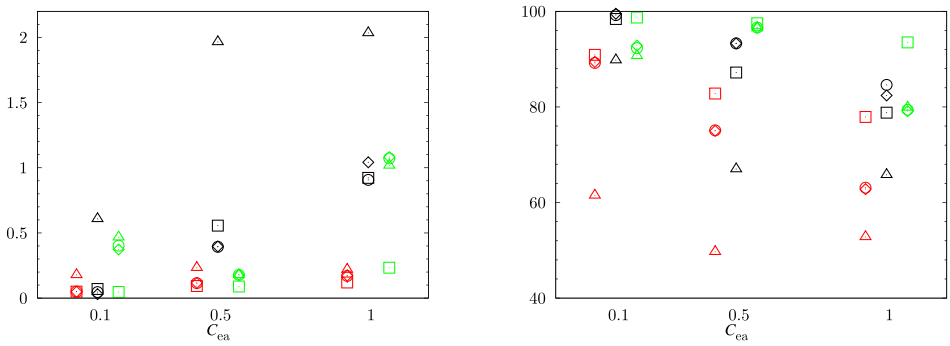

(a) PSO amplitude (%$\Delta\rho$) (b) PSO reduction (%)

Fig. 5. PSO of slow-moving shock: red symbols, M = 1.1; black symbols, M = 2.0; green symbols, M = 5.0; circles, $n = 5$, $N_e = 400$; squares, $n = 7$, $N_e = 400$; triangles, $n = 5$, $N_e = 200$; diamonds, $n = 5$, $N_e = 600$.

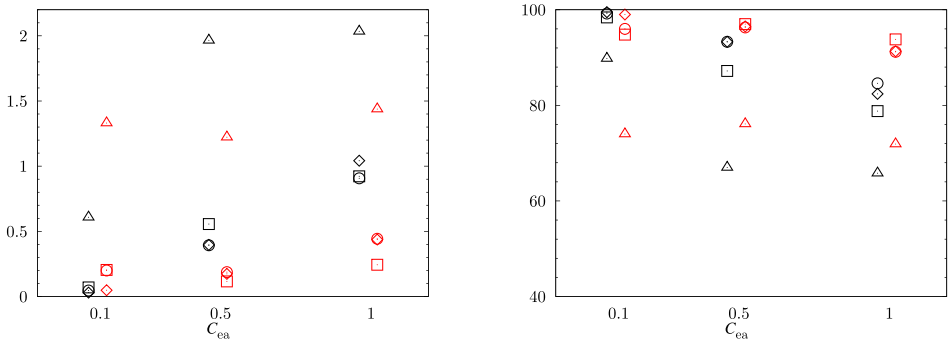

(a) PSO amplitude (%$\Delta\rho$) (b) PSO reduction (%)

Fig. 6. PSO of M = 2.0 slow-moving shock: black symbols, Laplacian AV; red symbols, physical AV; circles, $n = 5$, $N_e = 400$; squares, $n = 7$, $N_e = 400$; triangles, $n = 5$, $N_e = 200$; diamonds, $n = 5$, $N_e = 600$.

Overall, when dealing with very slow-moving shocks, the value of $C_{ea} = 0.5$ appears to be a fairly optimal choice for both the Laplacian and the physical formulations.

5. Evaluation of the shock velocity $\mathcal{D}$

One fundamental issue with the exponential averaging approach and the proposed calibration of the smoothing factor from Eq. (15) is the necessary knowledge of the shock displacement velocity $\mathcal{D}$ with respect to the mesh. In all the test presented in Section 4, the displacement speed of the shock was known in advance and hardcoded in the solver to compute the smoothing factor. On the other hand, in the vast majority of usual applications, this kind of information might not be available in advance. Moreover, the shock displacement speed might evolve during the computation. For instance, in simple problems of shock reflection, the incident and reflected shock speeds are different and the smoothing factor shall adapt to that change.

In this section, a methodology is proposed to evaluate the shock speed using either the density or the pressure profiles (or any other quantity that might identify the presence of a discontinuity).

Let a shock be propagating through a fluid with an absolute velocity $\mathbf{u}_a$. Notice that, if $\mathbf{u}$ is the fluid flow velocity, then the relative velocity of the shock with respect to the fluid can be computed as $\mathbf{u}_r = \mathbf{u}_a - \mathbf{u}$. Also note that the displacement speed $\mathcal{D}$ of the shock with respect to the mesh is related to the norm of $\mathbf{u}_a$. In other words, the computational mesh coincides with the absolute reference frame.

Let us identify, for instance, the shock front by the density profile. Because the profile propagates with the absolute velocity $\mathbf{u}_a$, the density must satisfy the advection equation

$$\frac{\partial \rho}{\partial t} + \mathbf{u}_a \cdot \nabla \rho = 0, \quad \Rightarrow \mathbf{u}_a \cdot \nabla \rho = -\frac{\partial \rho}{\partial t}. \tag{16}$$

Moreover, from the conservation of mass, the right-hand-side of Eq. (16) is:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0, \quad \Rightarrow -\frac{\partial \rho}{\partial t} = \nabla \cdot (\rho \mathbf{u}), \tag{17}$$

and hence, from Eqs. (16) and (17), the following identity is obtained:

$$\mathbf{u}_a \cdot \nabla \rho = \nabla \cdot (\rho \mathbf{u}). \tag{18}$$

By multiplying and dividing the left-hand-side by the norm of $\nabla\rho$, after the introduction of the unitary vector $\hat{\mathbf{n}} = \nabla\rho / \|\nabla\rho\|$ normal to the shock front (in the direction of the compressed gases), the above equation becomes

$$\mathbf{u}_a \cdot \frac{\nabla \rho}{\|\nabla \rho\|} \|\nabla \rho\| = (\mathbf{u}_a \cdot \hat{\mathbf{n}}) \|\nabla \rho\| = \nabla \cdot (\rho \mathbf{u}), \quad \Rightarrow \mathbf{u}_a \cdot \hat{\mathbf{n}} = \frac{\nabla \cdot (\rho \mathbf{u})}{\|\nabla \rho\|}, \tag{19}$$

which can be used to evaluate the norm of the absolute velocity (normal to the shock front), which is here identified with $\mathcal{D}$:

$$\mathcal{D} = \left| \mathbf{u}_a \cdot \hat{\mathbf{n}} \right| = \frac{|\nabla \cdot (\rho \mathbf{u})|}{\|\nabla \rho\|}. \tag{20}$$

By similar steps, in the case of inviscid flows, the absolute velocity of the shock can be obtained from the pressure using the pressure equation,

$$\frac{\partial p}{\partial t} + \mathbf{u} \cdot \nabla p + \rho a^2 (\nabla \cdot \mathbf{u}) = 0, \tag{21}$$

where a is the speed of sound. In this case, the shock speed becomes:

$$\mathcal{D} = \left| \mathbf{u}_a \cdot \hat{\mathbf{n}} \right| = \frac{1}{\|\nabla p\|} \left| \mathbf{u} \cdot \nabla p + \rho a^2 (\nabla \cdot \mathbf{u}) \right|, \tag{22}$$

where $\hat{\mathbf{n}} = \nabla p / \|\nabla p\|$. The advantage of the density is clearly in the simplicity of the relevant conservation law. Nonetheless, there might be cases in which the pressure signal could be a preferred choice.

Eq. (20) (resp. (22)) presents three major issues when used inside a high-order numerical solver: (a) the gradient of a high-order polynomial solution can be somewhat oscillatory, especially in the neighborhood of discontinuities; (b) the mass conservation equation (17) (resp. the pressure Eq. (21)) might not correspond *exactly* to what is numerically solved for (for example, when using a Laplacian AV, an additional diffusive term is present on the right-hand-side of Eq. (17)); (c) the division by the norm of the gradient of ρ (resp. p) is a quite bad-behaved operation, as this quantity can be zero or very close to zero in large regions of the flow.

Focusing for example on the density formulation, to overcome the first issue, the instantaneous value of the right-hand-side of Eq. (20) is integrated inside each element, rather than being computed at each solution point. More precisely, a gradient-weighted average of the shock speed is computed in every element as:

$$\mathcal{D}_e = \frac{\int_{V_e} \mathcal{D} \left\|\nabla\rho\right\| \mathrm{d}V}{\int_{V_e} \left\|\nabla\rho\right\| \mathrm{d}V} = \frac{\int_{V_e} \left|\nabla\cdot(\rho\mathbf{u})\right| \mathrm{d}V}{\left\|\nabla\rho\right\|_e}, \tag{23}$$

where V_e is the element volume and $\left\|\nabla\rho\right\|_e = \int_{V_e} \left\|\nabla\rho\right\| \mathrm{d}V$ has been introduced to simplify the notation hereafter. Notice that the integration is performed straightforwardly by Gauss quadrature at the solution points.

Concerning the second point, rather than using the theoretical transport equation for the selected variable, the actual residuals are adopted. In fact, the residuals automatically include all the viscous and/or AV terms, if any. Hence, using the residual of the mass conservation equation, $\mathcal{R}_\rho = \partial\rho/\partial t$, and following the same steps in Eqs. (16)–(20), the displacement speed becomes $\mathcal{D} = \left|\mathcal{R}_\rho\right| / \left\|\nabla\rho\right\|$. Should the pressure be used to detect the shock speed via Eq. (22), the displacement speed would be $\mathcal{D} = \left|\mathcal{R}_p\right| / \left\|\nabla p\right\|$, where the residual of the pressure is computed from the residuals $\mathcal{R}_\rho$, $\boldsymbol{\mathcal{R}}_{\rho\mathbf{u}}$ and $\mathcal{R}_{\rho E}$ of the mass, momentum and total energy equations, respectively:

$$\mathcal{R}_p = (\gamma-1)\left[\mathcal{R}_{\rho E} - \mathbf{u}\cdot\boldsymbol{\mathcal{R}}_{\rho\mathbf{u}} + \frac{1}{2}(\mathbf{u}\cdot\mathbf{u})\mathcal{R}_\rho\right]. \tag{24}$$

The above relation based on the residuals, has the practical advantage of being valid also in the viscous case without modification.

Eventually, Eq. (23) transforms into

$$\mathcal{D}_e = \frac{\int_{V_e} \left|\mathcal{R}_\phi\right| \mathrm{d}V}{\left\|\nabla\phi\right\|_e}, \tag{25}$$

where ϕ denotes, hereafter, either the density ρ or the pressure p.

Finally, concerning the third point, Eq. (25) is computed conditioned to a sufficiently large value of its denominator. Accordingly, if N_e is the number of elements and

$$\left\|\nabla\phi\right\|_{\max} = \max_{N_e}\left\{\left\|\nabla\phi\right\|_e\right\}, \tag{26}$$

then $\mathcal{D}_e$ becomes:

$$\mathcal{D}_e = \begin{cases} \dfrac{\int_{V_e} \left|\mathcal{R}_\phi\right| \mathrm{d}V}{\left\|\nabla\phi\right\|_e}, & \text{if } \left\|\nabla\phi\right\|_e > \beta_\nabla \left\|\nabla\phi\right\|_{\max}, \\ 0, & \text{otherwise,} \end{cases} \tag{27}$$

where the threshold coefficient β_∇ can be safely set in a relatively wide range of values without significantly affecting the results. A value of 0.01 will be adopted in all the tests presented hereafter.

5.1. Implementation

When a single shock is present in the domain, the shock speed from Eq. (27) can be used to calibrate the method globally, with the smoothing factor α determined, according to Eq. (15), using the maximum value of $\omega_e = \tau_{D,e}^{-1} = \mathcal{D}_e/\Delta x_e$ within regions where the shock is detected:

$$\omega_{\max} = \max_{N_e}\left\{\mathcal{D}_e/\Delta x_e \mid \varepsilon_e > \beta_\varepsilon\varepsilon_0\right\}, \tag{28}$$

where ε_e is the local element value of the AV, ε_0 is the relevant scaling factor from Eq. (8) and Δx_e is, for multi-dimensional problems, an average element size. The conditional coefficient β_ε is set equal to 0.01 hereafter.

When multiple shocks are present, a similar global approach can be used. In this case, however, considering that the most troublesome shock is the slowest, the smoothing factor might be better calibrated using the minimum non-zero value of ω within regions where the shock is detected:

$$\omega_{\min} = \min_{N_e}\left\{\mathcal{D}_e/\Delta x_e \mid \varepsilon_e > \beta_\varepsilon\varepsilon_0 \text{ and } \mathcal{D}_e/\Delta x_e > \epsilon\right\}, \tag{29}$$

where ϵ is a small threshold value of the order of the machine precision, which is introduced to prevent a zero value of $\omega_{\min}$.

Another possibility is to calibrate the global smoothing factor using a suitable average value. For instance, to restrict the averaging process to the regions where shocks are present, the average, weighted by the AV, can be used:

$$\omega_{\text{avg}} = \frac{\sum_{e=1}^{N_e} \varepsilon_e \mathcal{D}_e/\Delta x_e}{\sum_{e=1}^{N_e} \varepsilon_e}. \tag{30}$$

The problem with such global approaches based on one unique value of ω is that the resulting smoothing factor would be optimal for one shock at most. In particular, using the maximum values would result optimal for the fastest shock in the domain. As far as the slower shocks are concerned, these will benefit some PSO reduction but the reduction would be sub-optimal. On the other hand, the global calibration based on the minimum shock speed would introduce an excessive time delay on the faster shocks. From Fig. 4(a), it is clear that an excessive time delay can result in an insufficient amount of AV being injected in the neighborhood of the faster shocks and, as a result, these might suffer from excessive Gibbs phenomena or even instability. At the same time, the AV space localization would be penalized and excessive dissipation would result in the neighborhood of the fast shocks. Finally, a global approach based on the average would mitigate the above-mentioned problems for the faster shocks but would be sub-optimal overall for the slower shocks.

It is clear that the best approach would be to identify and follow the different shocks in the domain and calibrate the exponential smoothing factor in the neighborhood of every shock individually. To this end, the approach proposed here utilizes a tool borrowed from *machine learning* techniques to discriminate between regions of different shock speed and determine suitable global means. The calibration of the exponential smoothing is then performed in the neighborhood of every shock depending on the global mean which is closest to its displacement speed.

As a first step, the ensemble of all the *meaningful* values of ω within the whole domain is collected. In practice, this ensemble contains the values of ω measured where gradients are sufficiently high and the AV is non-zero:

$$\mathcal{X} = \left\{x_i \equiv \omega_e \mid \left\|\nabla\phi\right\|_e > \beta_\nabla \left\|\nabla\phi\right\|_{\max} \text{ and } \varepsilon_e > \beta_\varepsilon\varepsilon_0\right\}, \tag{31}$$

where $\beta_\nabla = \beta_\varepsilon = 0.01$ and

$$\omega_e = \frac{\mathcal{D}_e}{\Delta x_e} = \frac{\int_{V_e} \left|\mathcal{R}_\phi\right| \mathrm{d}V}{\left\|\nabla\phi\right\|_e \Delta x_e}. \tag{32}$$

Supposing that K shocks are present in the flow, it is reasonable to expect that the values within $\mathcal{X}$ will be *clustered* around a number of typical values ϑ_j, one for each shock in the domain. To find these typical values—or optimal cluster centroids—the power k-means clustering technique is adopted [28]. Among the many different k-means algorithms, this particular one has very low sensitivity on the cluster initiation while maintaining a relatively low algorithmic complexity.

Given a set of initial centroids $\vartheta_{0,j}$, with $j = 1, \ldots, K$, the power k-means clustering procedure involves the steps described in Algorithm

Algorithm 1 power k-means clustering

1: Initialize $s_0 > 0$ and centroids $\vartheta_{0,j}$ with $j = 1, \dots, K$
2: Initialize $f_0 \leftarrow \sum_{i=1}^{|\mathcal{X}|} \left[\frac{1}{K} \sum_{l=1}^{K} (x_i - \vartheta_{0,l})^{-2s_0} \right]^{-1/s_0}$
3: **while** %ch > 0.01 and $m \leq m_{\max}$ **do**
4: **for all** points x_i in $\mathcal{X}$ **do**
5: $w_{m,ij} \leftarrow \left[\sum_{l=1}^{K} (x_i - \vartheta_{m,l})^{-2s_m} \right]^{-(1+1/s_m)} (x_i - \vartheta_{m,j})^{-2(s_m+1)}$
6: **end for**
7: $\vartheta_{m+1,j} \leftarrow \left(\sum_{i=1}^{|\mathcal{X}|} w_{m,ij} \right)^{-1} \sum_{i=1}^{|\mathcal{X}|} w_{m,ij} x_i$
8: $f_{m+1} \leftarrow \sum_{i=1}^{|\mathcal{X}|} \left[\frac{1}{K} \sum_{l=1}^{K} (x_i - \vartheta_{m+1,l})^{-2s_m} \right]^{-1/s_m}$
9: %ch $\leftarrow |f_{m+1} - f_m| / f_m$
10: $s_{m+1} \leftarrow \eta s_m$
11: **end while**

1, where m denotes the iteration counter.[3] Notice that the convergence is checked on the performance function $f_m = f(\vartheta_{m,j})$ and, in particular, the iteration is stopped when the percentage change of f_m falls below 1% of the value at the previous iteration. Optionally, it might be useful to check the standard deviation of the data in $\mathcal{X}$ beforehand and skip the power k-means clustering approach if it is sufficiently small compared to the average value ω_{avg} from Eq. (30) (say 1% ω_{avg}).

Once the typical values ϑ_j are obtained for ω, to calibrate the exponential smoothing locally for every shock, the method proposed is based on the evaluation of an element based value of ω_e satisfying the following properties:

(a) ω_e shall be equal to the closest typical value in the neighborhood of every shock;
(b) ω_e shall have a non-zero, local, lower bound to avoid the AV to become *frozen*[4] in regions where steady shocks are present;
(c) ω_e shall automatically return to a sufficiently high maximum value where no shocks are present, such that, in these regions, the AV promptly recovers a zero value.

In order to satisfy these properties the element-wise value of ω_e^i, at the ith iteration is computed as:

$$\omega_e^i = \begin{cases} \dfrac{\int_{V_e} \left| \mathcal{R}_\phi \right| \mathrm{d}V}{\|\nabla\phi\|_e \, \Delta x_e}, & \text{if } \|\nabla\phi\|_e > \beta_\nabla \|\nabla\phi\|_{\max}, \\ \vartheta_{j,\text{cl}}, & \text{if } \|\nabla\phi\|_e \leq \beta_\nabla \|\nabla\phi\|_{\max} \text{ and } \varepsilon_e > \beta_\varepsilon \varepsilon_0, \\ \dfrac{(u+a)_{\max}}{\Delta x_e}, & \text{otherwise,} \end{cases} \tag{34}$$

where $\vartheta_{j,\text{cl}}$ is the closest value to ω_e^{i-1}, from the previous iteration, among the typical values ϑ_j. To initialize the simulation, an initial value is set, everywhere, using the maximum speed of sound: $\omega_e^0 = a_{\max}/\Delta x_e$. According to Eq. (34), ω_e^i will assume the locally computed value $D_e/\Delta x_e$ in regions where the density gradient is sufficiently high. In the neighborhood of the shocks, where the AV is active, ω_e^i will remain equal to the closest typical value—or centroid—of the cluster the shocks belongs to. In all the other regions, where no shocks are detected, the exponential smoothing is set with the time-scale of the fastest wave speed to ensure fast recovery of the AV.

Finally, the exponential smoothing factor for the AV is obtained, at every iteration, from Eq. (15):

$$\alpha_e^i = 1 - \exp\left(-C_{\text{ea}} \frac{\omega_e^i \Delta t}{n_{\text{RK}}} \right). \tag{35}$$

The actual implementation is summarized in Algorithm 2 and will be referred to as *adaptive exponential smoothing* hereafter.

Algorithm 2 Assign ω_e^i at i-th iteration and set α_e^i

1: Initialize $\omega_e^0 = a_{\max}/\Delta x_e$ and $\beta_\nabla = \beta_\varepsilon = 0.01$
2: **for all** element e **do**
3: **if** $\|\nabla\phi\|_e > \beta_\nabla \|\nabla\phi\|_{\max}$ **then**
4: $\omega_e^i \leftarrow \int_{V_e} \left| \mathcal{R}_\phi \right| \mathrm{d}V / \left(\|\nabla\phi\|_e \, \Delta x_e \right)$
5: **else if** $\varepsilon_e > \beta_\varepsilon \varepsilon_0$ **then**
6: $d_j \leftarrow \left(\omega_e^{i-1} - \vartheta_j \right)^2$
7: $d_{\min} \leftarrow \min_j \{ d_j \}$
8: $k \leftarrow$ index of $d_{\min}$
9: $\omega_e^i \leftarrow \vartheta_k$
10: **else**
11: $\omega_e^i \leftarrow (u+a)_{\max}/\Delta x_e$
12: **end if**
13: $\alpha_e^i \leftarrow 1 - \exp\left(-C_{\text{ea}} \omega_e^i \Delta t / n_{\text{RK}} \right)$
14: **end for**

Table 4
Shock speed evaluation accuracy ($C_{\text{ea}} = 0.5$).

	Exact D	Eq. (27)	
		$\phi \equiv \rho$	$\phi \equiv p$
Test 1.	0.118	0.118 ± 0.0016	0.118 ± 0.0018
Test 2.	3.550	3.542 ± 0.0074	3.543 ± 0.0140
	1.446	1.437 ± 0.0082	1.441 ± 0.0095
Test 3.	3.313	3.305 ± 0.0074	3.307 ± 0.0137
	2.366	2.365 ± 0.0044	2.366 ± 0.0021
	1.999	1.992 ± 0.0062	1.992 ± 0.0140
	0.599	0.597 ± 0.0029	0.599 ± 0.0032

Table 5
PSO amplitude with power k-means adaptive exponential smoothing ($C_{\text{ea}} = 0.5$).

	D	K	PSO amplitude ($\%\Delta\rho$)		PSO reduction (%)	
			Lapl. AV	Phys. AV	Lapl. AV	Phys. AV
Test 1.	0.118	1/2	0.381	–	93.5	–
Test 2.	1.446	1/2	0.043	0.046	93.7	98.3
Test 3.	1.999	2/3	~0.06	0.198	~93	95.3
	0.599	2/3	0.137	0.256	92.6	95.3

5.2. Numerical tests

The shock velocity evaluation method is tested on three one-dimensional tests:

1. A Mach 2 moving shock at 5% its displacement speed in fluid at rest. This is the same test described in Section 3 and the expected shock speed is equal to 0.118 (computational units).
2. Two colliding Mach 3 shocks (equivalent to a shock reflection problem). In this case, setting to 1 the pressure and density in the fluid at rest between the shocks and the corresponding values upstream of the shocks according to Rankine–Hugoniot conditions for the selected Mach number, the incident and reflected shocks will have $D = 3.550$ and 1.446 (Mach 2.1), respectively.
3. Two colliding shocks at Mach 2 and 2.8, respectively. Here two different values of the shock speed are simultaneously present.

[3] Some extra care shall be taken when dealing with steps 2, 5 and 8 of Algorithm 1, as the relevant summations involve potential divisions by zero. The problem can be easily overcome by using the identity [29]

$$\left[\sum_i d_i^{-2s} \right]^{-p} = \left(d_{\min}^2 \right)^{sp} \left[1 + \sum_{i \neq \min} \left(\frac{d_{\min}^2}{d_i^2} \right)^s \right]^{-p}, \tag{33}$$

which does not pose any problem even when $d_{\min}^2 = \min\{d_i^2\} = 0$.

[4] From Eq. (15), $\alpha_e = 1 - \exp(-C_{\text{ea}} \omega_e \Delta t / n_{\text{RK}})$ and $\alpha_e = 0$ for $\omega_e = 0$. A zero value of the smoothing factor implies that the exponential averaging does not weight present realizations (*cf.* Eq. (10)) and, therefore, the AV cannot change anymore.

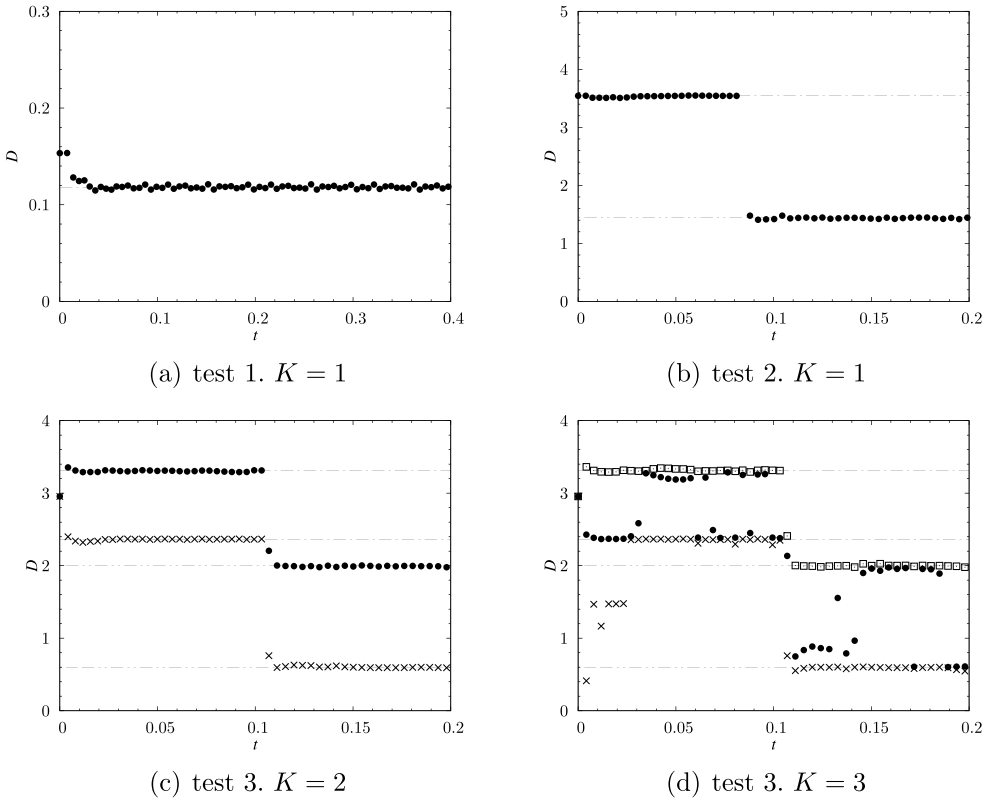

Fig. 7. Shock speed evaluation by the power k-means clustering, using the density formula (*cf.* Eq. (27)): filled circles, computed shock speed (1st cluster centroid); crosses, computed shock speed (2nd cluster centroid); open squares, computed shock speed (3rd cluster centroid); dash-dotted line, theoretical (incident) shock speed; dash–dot-dotted line theoretical (reflected) shock speed.

For unitary pressure and density in the fluid at rest between the shocks, the incident shocks have $D = 2.366$ and 3.313, respectively. After collision, two reflected shocks and a contact discontinuity between them are present. The reflected shocks have $D = 1.999$ (Mach 2.26) and 0.599 (Mach 1.62), respectively, whereas the contact discontinuity has a displacement velocity of 0.853.

The first two tests are designed to check the accuracy of the proposed shock speed detection method and its capacity to adapt to shock speed changes, whereas the last test is designed to check the ability of the power k-means clustering technique to find the correct clusters in the data and the relevant centroids, *i.e.*, the propagation velocities of the shocks. In the first two tests, the number of clusters K is set equal to 1 or 2, whereas, for the third test, $K = 2$ or 3.

The time history of the computed shock speed using power k-means clustering via Eq. (27), with $\phi \equiv \rho$, is depicted in Fig. 7 for the three test cases. The accuracy of the results obtained using both the density or the pressure in Eq. (27) (mean value and standard deviation) is reported in Table 4. Overall, both formula produce sufficiently accurate estimations of D and the power k-means algorithm is capable to identify shocks and adapt to changes in their speed. Some oscillations in the computed shock speed are visible, but these do not seem to negatively impact the results in terms of PSO (see below). Concerning the parameter K of the power k-means clustering procedure, it is worthwhile noticing that selecting K greater than the actual number of different shocks in the flow results in the k-means algorithm to find additional cluster centroids. However, these are close to the centroids of the existing shocks. This can be seen in Fig. 7(d), relevant to the third test executed with $K = 3$. As it can be seen, the third additional shock speed (or cluster centroid) oscillates between the typical expected values characterizing the two shocks in the flow. Therefore, given that Eq. (34) selects the closest centroid, the exponential averaging calibration remains unaffected by the selected value of K (provided, of course, that K is high enough to account for all the different shocks in the flow).

A summary of the measured PSO amplitudes and reductions for the different test cases is reported in Table 5. The data include results from both the Laplacian and the physical AV models. As anticipated, the impact of the selected value of K is, at most, marginal and, for each test, the measured PSO are unaffected by K. Concerning the first test case, in particular, the adaptive power k-means calibration performs the same as the calibration based on the *a priori* known shock speed, with PSO amplitudes of about 0.4 $\%\Delta\rho$ and PSO reductions of about 93% for both approaches (*cf.* Tables 2 and 5).

The density and AV profiles for the three test cases are depicted in Fig. 8, where the results without exponential smoothing are also shown for comparison. Notice that, the shown results for the second and third tests have been obtained using the physical AV. This choice, in fact, is more suitable for shock collision problems, which can be particularly prone to produce overheating errors. The results obtained with the Laplacian AV model (not shown) are anyway comparable, except for some localized disturbances due to the presence of an artificial dissipative term in the density equation and the lack of strict physical consistency between momentum and energy transport (see Ref. [5] for a detailed discussion).

6. Results and discussion

In the present Section, the results obtained using the exponential averaging technique are discussed on some selected one- and two-dimensional test cases. Comparisons will be presented with similar computations performed without exponential averaging. It shall be noticed that, unless stated otherwise, the baseline numerical setup uses the Roe flux with entropy fix [21,22] for the inviscid fluxes, whereas the interior penalty (IP) flux (with $\alpha_{\text{IP}} = 0.5$) is adopted for the viscous fluxes. In the case of inviscid computations, the molecular diffusion

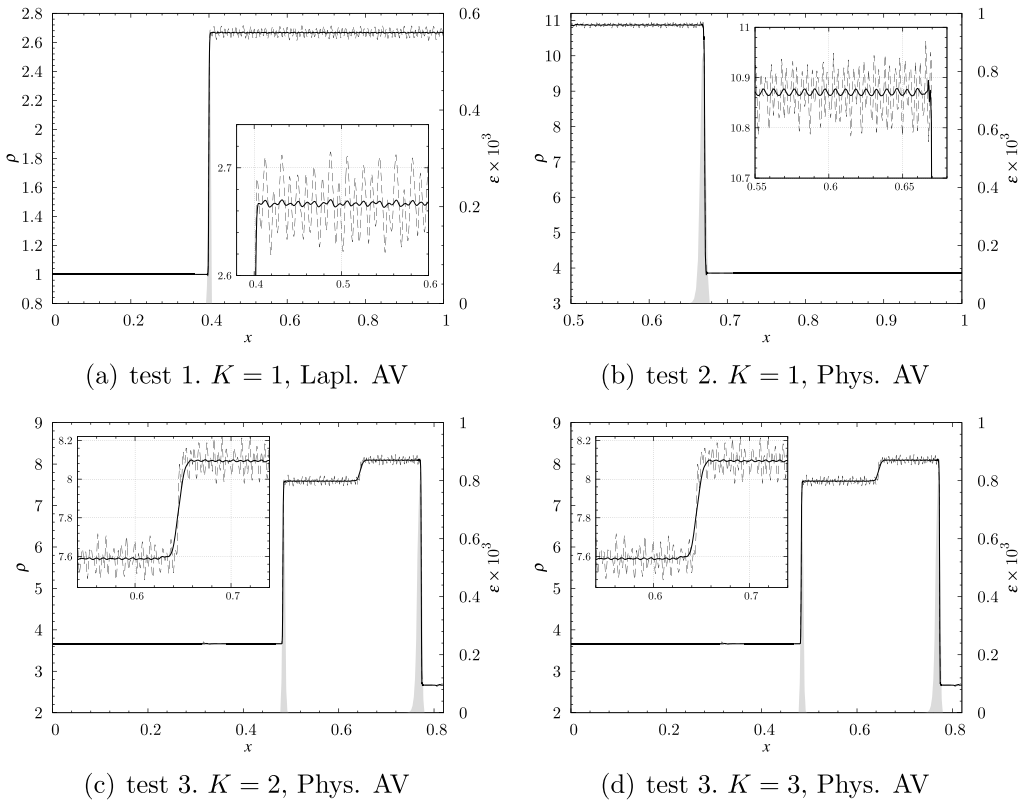

Fig. 8. Comparison of density profile with and without power k-means adaptive exponential smoothing: solid line, density (left axis); gray filled line, shock-capturing viscosity (right axis); thin dashed line, density without exponential smoothing.

is switched off and the IP flux is used for the artificial viscosity term only. Time integration is performed with the RK45-SSP Runge–Kutta scheme [23] with 0.4 CFL and Fourier coefficients. Shock-capturing is enforced by means of a Laplacian or physical term, whose viscosity is scaled, with the parameter $C_\varepsilon = 1$, on the maximum wave speed within the computational domain. The positivity-preserving procedure [26] for the density and pressure will also be activated whenever needed.

6.1. Sod's shock tube

The Sod tube Riemann problem [27] is considered in this Section. On this particular test, oscillations in the density of about 0.4–0.5% of the density jump across the contact discontinuity where observed in Ref. [4], regardless of the order of accuracy. Numerical tests are performed by integration of the Euler equations on either 200 or 600 uniform elements, using, respectively 10th- and 8th-order SD discretizations. The flow is initialized at $t = 0$ setting $(\rho, u, p) = (1, 0, 1)$ and $(\rho, u, p) = (0.125, 0, 0.1)$ in the left and right halves of the domain, respectively. Then, the solution is checked at $t = 0.2$, when a rarefaction wave, a contact discontinuity and a shock are established.

The relevant results, obtained using both the Laplacian and the physical AV models, are reported in Fig. 9. As it can be seen, depending on the grid resolution, the order and the AV model, the computations without exponential smoothing are characterized by PSO with amplitudes as high as about 1.2% of the density jump. Notice that the oscillations to the left of the contact discontinuity originate behind the shock. This can be readily verified by inspection of the pressure profile (not shown). The PSO reductions obtained thanks to the adaptive exponential smoothing technique are quite significant and range between 96.1 and 99.5%.

6.2. Shu–Osher shock tube

The Shu Osher shock tube problem [30] involves a Mach 3 shock, traveling at its maximum speed in fluid at rest ($D^* = 3.549648$), interacting with a sinusoidal wave in density. For such a propagation velocity, no PSO are observed[5] and, in principle, the exponential smoothing approach is not required. This notwithstanding, the test is performed to assess whether or not the penalty in AV localization promoted by the exponential smoothing translates into a significant increase in artificial dissipation.

The test is performed by integration of the Euler equations on 400 uniform elements using a 5th-order SD discretization over a domain of unitary length. The flow is initialized at $t = 0$ setting:

$$\begin{aligned} &\rho = 3.857143, && u = 2.629369, && p = 10.33333, && \text{for } x \le 0.1, \\ &\rho = 1 + 0.2\sin[50(x - 0.5)], && u = 0.0, && p = 1.0, && \text{for } x > 0.1. \end{aligned}$$

The density profile at $t = 0.2$ is shown in Fig. 10(a), where the solutions obtained with and without exponential smoothing are compared. Some differences between the two computations are visible only in the fine structure in the density profile (*cf.* closeup view). In particular, the expected reduction in AV localization brought by the exponential smoothing procedure promotes a slight increase in overall dissipation. As a result, the fine density structures appear slightly *damped* with respect to the computation without exponential smoothing.

Of course, such a marginal loss of detail shall be put in perspective of applications where severe PSO are present and can, potentially, have a significant impact on the solution. To this end, the test can be modified such as to have the same Mach 3 shock propagating with a

[5] The measured amplitude of PSO for an identical moving shock, in the absence of interaction with the sinusoidal density wave, was down to machine precision.

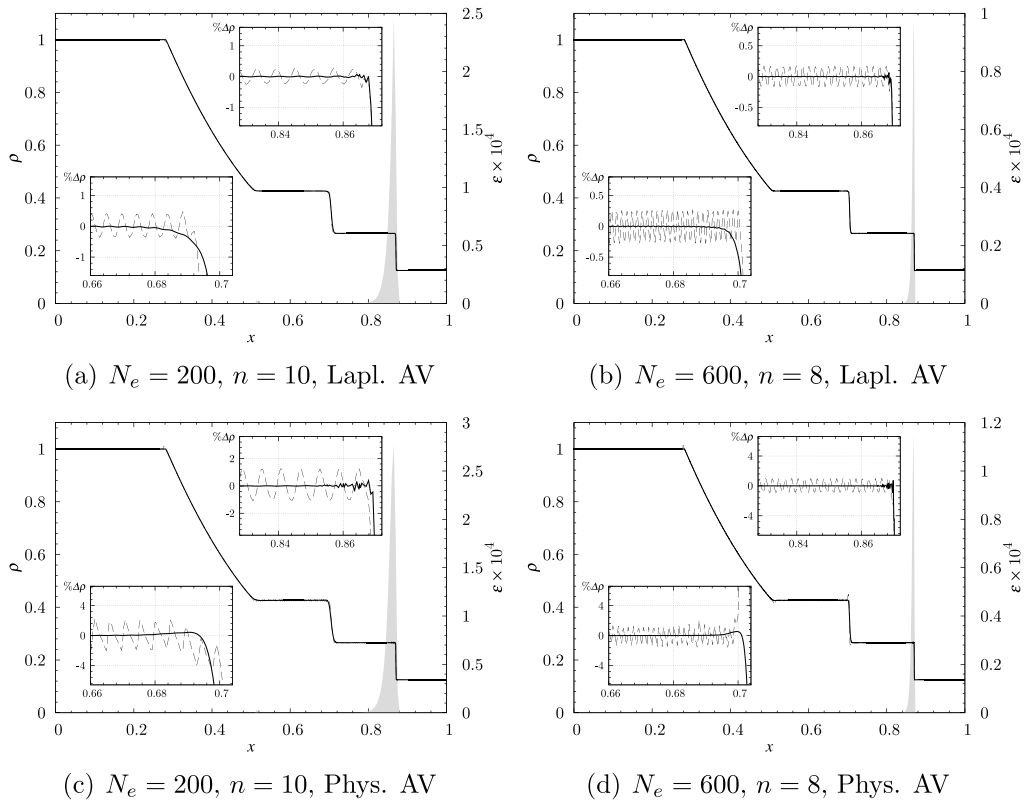

Fig. 9. Sod tube Riemann problem: solid line, density (left axis); gray filled line, shock-capturing viscosity (right axis); thin dashed line, density without exponential smoothing.

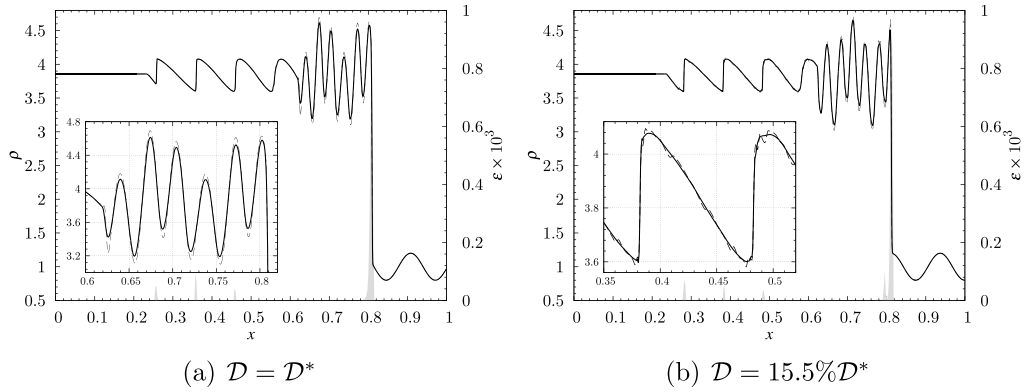

Fig. 10. Shu–Osher shock tube problem: solid line, density (left axis); gray filled line, shock-capturing viscosity (right axis); thin dashed line, density without exponential smoothing.

much slower velocity with respect to the mesh. The equations are then solved in a moving reference frame with some extra care in properly setting boundary conditions. For instance, moving the reference frame with a velocity $U_{\text{ref}} = 3.0$, the shock appears propagating with a relative speed of about 15.5% of its maximum speed in fluid at rest $\mathcal{D}^*$. The shock is, in this case, initialized at $x = 0.7$ and the solution checked at $t = 0.2$. The results of such a *slow-moving adaptation* of the Shu Osher shock tube problem are shown in Fig. 10(b). As it can be clearly seen in the closeup view of the train of waves to the left of the fine structure in the density, the shock now produce significant PSO, which perturb the flow downstream. Using the adaptive exponential smoothing ($K = 1$), on the other hand, the PSO are almost completely suppressed. Indeed, this kind of scenario represents a very good example in which the gain in term of PSO suppression is well worth the marginal penalty in AV localization and overall dissipation.

Table 6
Two-dimensional Riemann problems.

Case	Quadrant	ρ	u	v	p	t_f
4	1	1.1	0.0	0.0	1.1	0.25
	2	0.5065	0.8939	0.0	0.35	
	3	1.1	0.8939	0.8939	1.1	
	4	0.5065	0.0	0.8939	0.35	
12	1	0.5313	0.0	0.0	0.4	0.25
	2	1.0	0.7276	0.0	1.0	
	3	0.8	0.0	0.0	1.0	
	4	1.0	0.0	0.7276	1.0	

6.3. Two-dimensional Riemann problems

In this section, results on two selected two-dimensional Riemann problems are analyzed, namely the case 4 and the case 12 reported

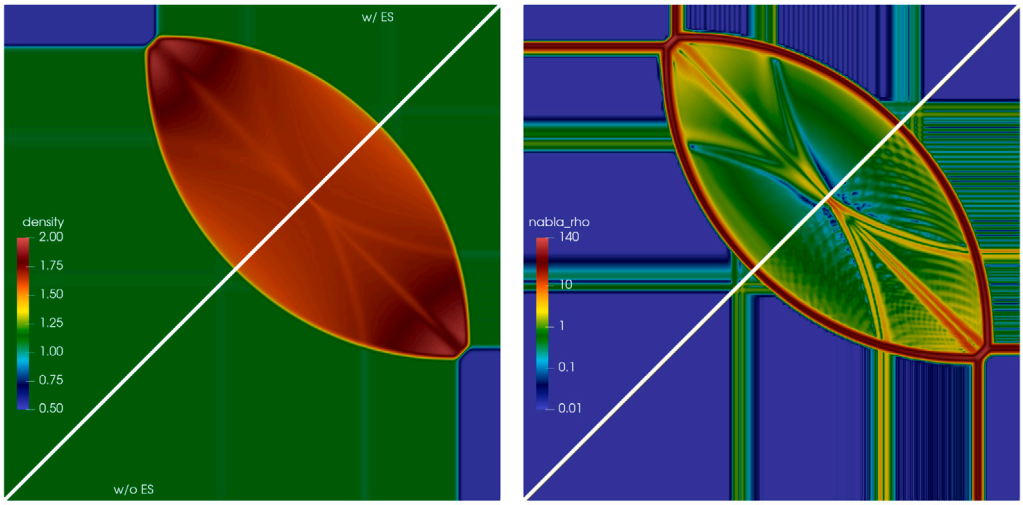

(a) 128×128, $n = 4$ (b) 128×128, $n = 4$

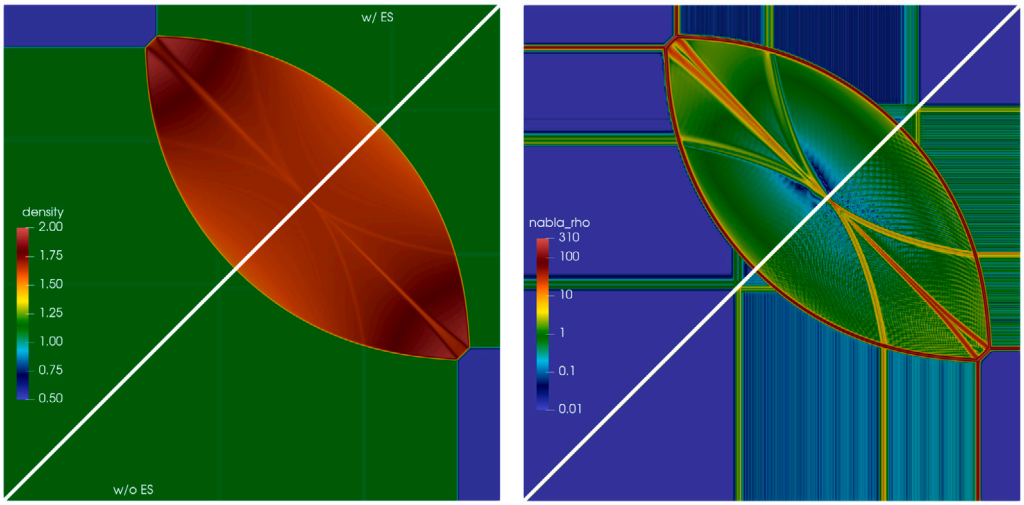

(c) 256×256, $n = 5$ (d) 256×256, $n = 5$

Fig. 11. Contours of density (a, c) and norm of density gradient (b, d) for the case 4 test with (top-left half) and without (right-bottom half) adaptive exponential smoothing.

in [31,32]. The computational domain is a square of unitary side length discretized with either 128×128 4th-order quadrilateral elements or 256×256 5th-order quadrilateral elements. The total number of degrees of freedom (DoF) is 262 144 and 1 638 400, respectively. Case 4 is computed on both meshes, whereas case 12 is computed on the higher resolution only.[6]

The domain is subdivided in four identical quadrants and the Riemann problem is defined by initial constant states in each quadrant (numbered counterclockwise, from 1 to 4, starting from the upper-right quadrant). The initial states and the final time t_f for the presented tests are reported in Table 6.

The computations have been performed with and without the adaptive exponential smoothing. Concerning the power k-means procedure, the cluster number K has been set equal to 2. Time integration has been performed with the RK45-SSP, with 0.5 CFL and Fourier coefficients.

With regards to case 4, Fig. 11 shows the contours of density and the norm of its gradient. Fig. 12 shows density and AV profiles along a horizontal line at $y = 0.97$—which crosses the main shock located in the 2nd quadrant at the final time—and along a 45° line, respectively.

On both mesh resolutions, from the density contours, an increase in the level of dissipation produced by the exponential smoothing is visible within the high-density region in the upper-right part of the domain. This is particularly evident on the slip lines (oriented at −45°), issuing from the two pairs of triple points located in the 2nd and 4th quadrants, and is readily confirmed by the AV profile along the 45° line [*cf.* Fig. 12(b) and (d)], where a significant penalty in AV localization is observed for the coarsest mesh. Not necessarily a drawback of the proposed method, other structures which appear smoother (weaker) when the exponential smoothing is activated are the spurious waves—due to the shocks being initialized as sharp profiles—which trail behind the leading shocks. Increasing the resolution, the observed AV de-localization is reduced, though, and the difference in artificial dissipation becomes less evident.

As it can be clearly seen in the contours of density gradient, Fig. 11(b) and (d), the leading shocks located in the 2nd and 4th quadrants are affected by relatively strong PSO when the exponential smoothing is deactivated. These oscillations are strong enough to penetrate and perturb the above mentioned high density region. This is

[6] It shall be noted that, for case 4, given the particular initial state in the 2nd, 3rd and 4th quadrants, and in order to allow the development of the flow in the 2nd and 4th quadrants without perturbations from the boundary, the boundary condition at the left and bottom boundaries has been enforced using self-periodic boundary elements as discussed in more detail in Section 6.5 for the double Mach reflection problem.

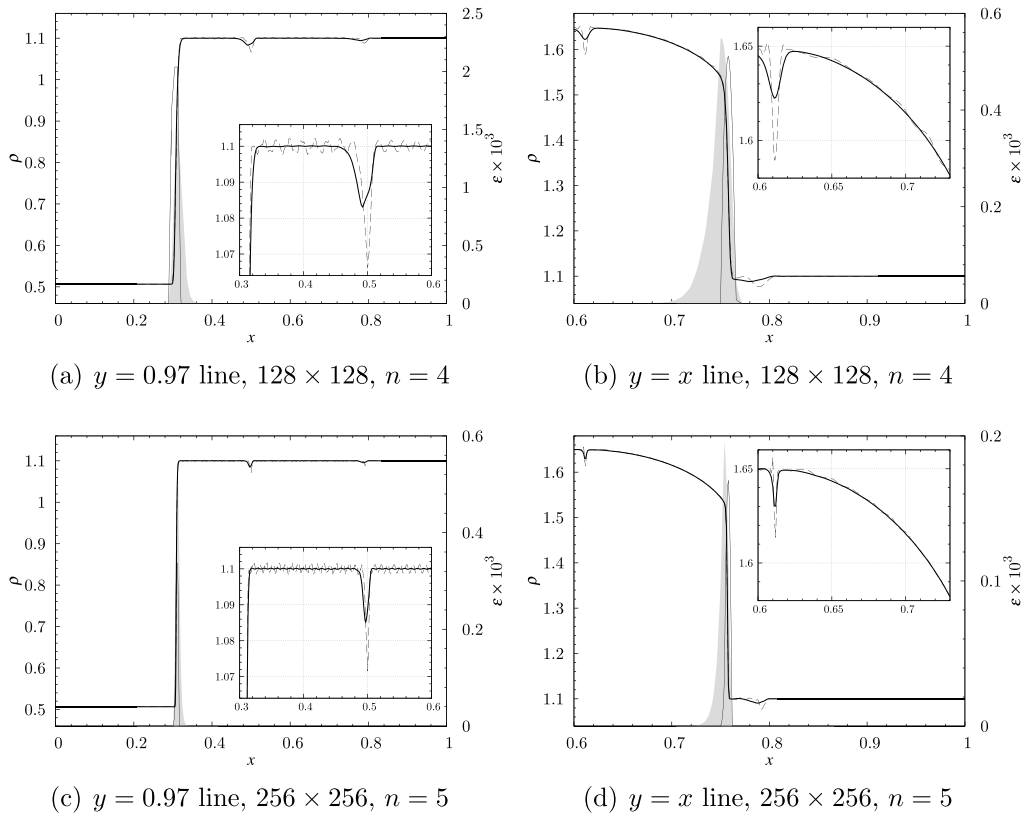

Fig. 12. Density and AV profiles for case 4: solid line, density with adaptive exponential smoothing, $K = 2$ (left axis); thin dashed line, density without exponential smoothing (left axis); gray filled line, shock-capturing viscosity with adaptive exponential smoothing (right axis); thin solid line, shock-capturing viscosity without exponential smoothing (right axis).

readily confirmed by looking at the density profiles in Fig. 12. The use of exponential smoothing promotes the complete suppression of these oscillations and, despite the reduction in AV localization, no significant reduction in the sharpness of the shocks is observed.

Switching to case 12, Fig. 13 shows the contours of density, the norm of its gradient and AV, whereas Fig. 14 shows the density and AV profiles along a horizontal line $y = 0.95$—which crosses the main shock located in the 1st quadrant at the final time—and along a 45° line, respectively. The main shock, in particular, is characterized by PSO in density of about 0.1% of the density jump. These are almost completely suppressed by the adaptive exponential smoothing [*cf.* closeup inset in Fig. 14(a)].

As it can be seen in Figs. 13 and 14, the AV localization, again, is somewhat reduced and the penalty in terms of AV localization is particularly accentuated along the 45° direction, at the location where the different shocks join in triple points. Analogously to the previous case, the 45° oriented slip lines, issuing from the pair of triple points in the 1st quadrant, appear slightly smoother when the exponential smoothing is active.

A notable effect of the exponential smoothing concerns the contact discontinuities in the 3rd quadrant of case 12. In fact, the reference simulation performed without exponential smoothing presents spurious (and most probably unsteady) peaks of AV at the boundary, between the 2nd and 3rd quadrants and between the 3rd and 4th quadrants. One of these perturbations is visible as a red spot of AV at the bottom boundary in Fig. 13(c). These peaks promote the destabilization of the contact discontinuities which originate at those points, as it can be clearly observed in Fig. 13(a) and (b). Such destabilization is of course non-physical [31,32] and is completely avoided thanks to the time delay introduced in the AV by the exponential smoothing.

As in the previous example, spurious waves due to the shocks being initialized as sharp profiles are made weaker by the exponential smoothing. Other spurious waves, most probably caused by the above-mentioned peaks of AV at the boundary, are not present in the simulation featuring exponential smoothing [*cf.* 3rd and 4th quadrants in Fig. 13(b)]. These, of course, can be considered as desirable side effects of the exponential smoothing.

6.4. Shock/vortex interaction

The interaction between a two-dimensional inviscid vortex and a $\mathrm{M}_s = 3$, slow-moving (the shock speed is set at 5% of its propagation speed in fluid at rest D^*), normal shock is here analyzed. The square computational domain of dimensions $H \times H$ consists of 128×128 quadrilateral elements. Setting the order n equal to 5, the total number of DoF is 409 600. Inflow (prescribed density and velocity) and outflow (prescribed pressure) conditions are set at the left and the right boundaries, respectively, whereas periodic conditions are used at the top and bottom boundaries.

The shock is first initialized according to the Rankine–Hugoniot jump conditions such as to obtain a Mach 3 stationary shock. Hence the left and right velocities are modified to account for a reference frame moving with a translation velocity equal to $D = 5\% D^*$:

$$u_{L,\mathrm{rel.}} = u_L - D, \quad u_{R,\mathrm{rel.}} = u_R - D, \quad \text{with} \quad D = 5\% D^* = 5\% u_L, \tag{36}$$

where the L and R subscripts denote the regions upstream and downstream of the shock, respectively.

Upstream of the shock, in the middle of the region to its left, a homentropic Taylor vortex [33] is initialized according to the following relations:

$$\left[\frac{\rho(\mathbf{x})}{\rho_L}\right]^{\gamma-1} = \left[\frac{p(\mathbf{x})}{p_L}\right]^{\frac{\gamma-1}{\gamma}} = 1 - \frac{\gamma-1}{2}\mathrm{M}_v^2 \exp(1-\xi_v^2-\eta_v^2), \tag{37}$$

$$\frac{u_{L,\mathrm{rel.}} - u(\mathbf{x})}{|u_L|\,\eta_v} = \frac{v(\mathbf{x})}{|u_L|\,\xi_v} = \frac{\mathrm{M}_v}{\mathrm{M}_s} \exp\left(\frac{1-\xi_v^2-\eta_v^2}{2}\right), \tag{38}$$

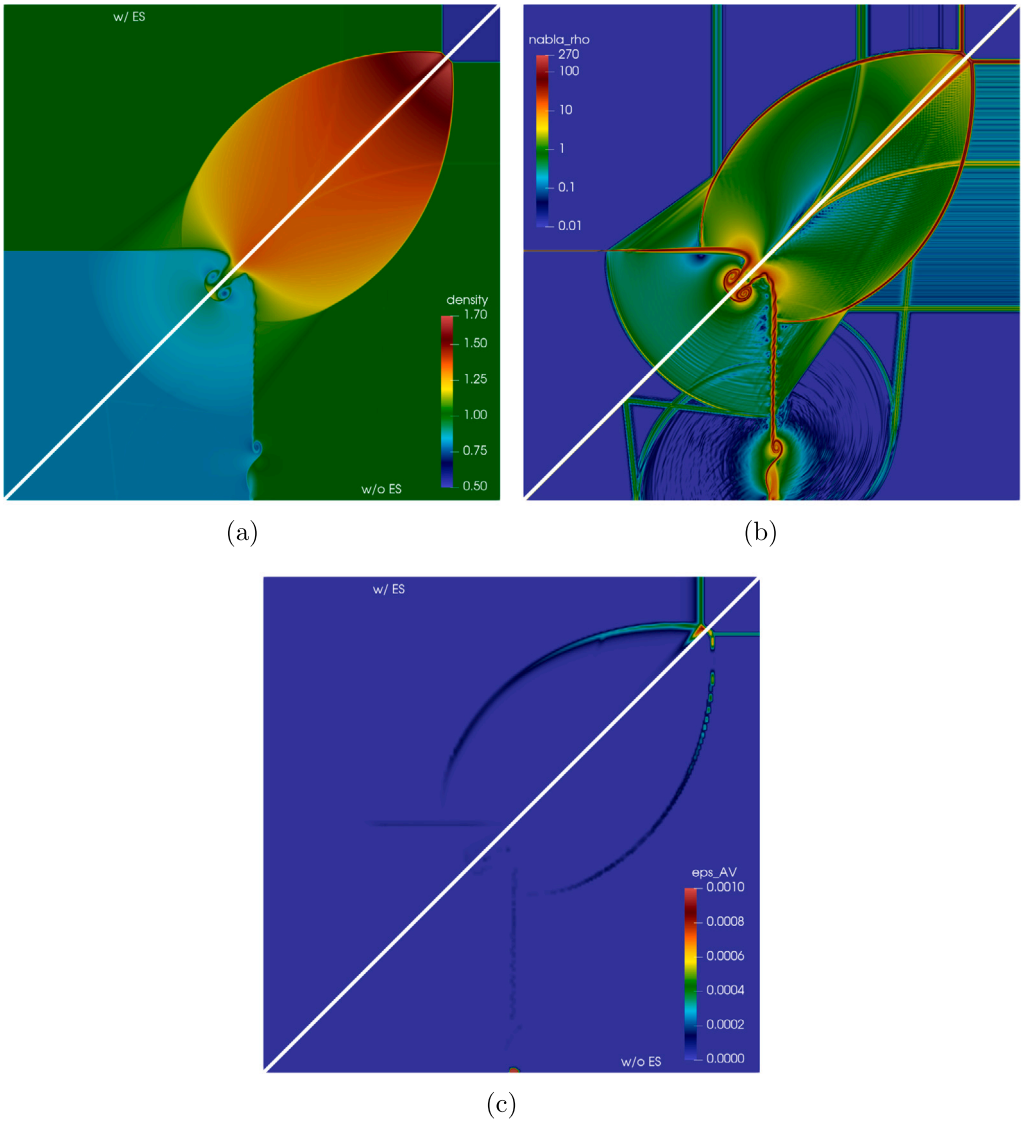

Fig. 13. Contours of density (a), norm of density gradient (b) and AV (c) for the case 12 test with (top-left half) and without (right-bottom half) adaptive exponential smoothing.

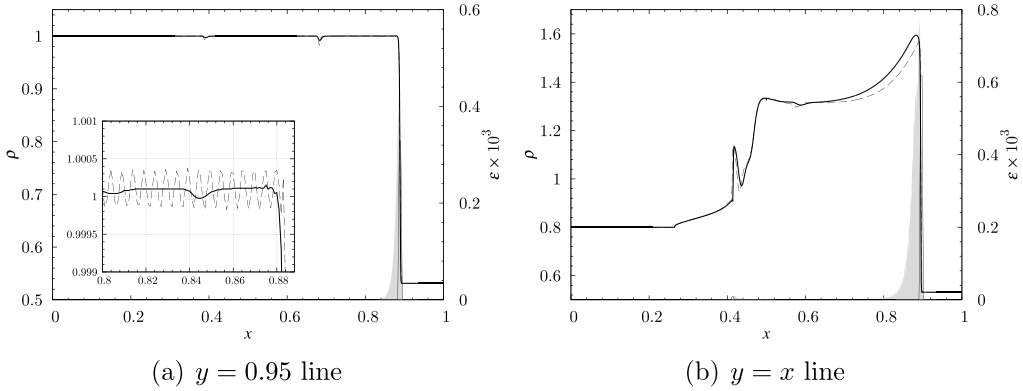

Fig. 14. Density and AV profiles for case 12: solid line, density with adaptive exponential smoothing, $K = 2$ (left axis); thin dashed line, density without exponential smoothing (left axis); gray filled line, shock-capturing viscosity with adaptive exponential smoothing (right axis); thin solid line, shock-capturing viscosity without exponential smoothing (right axis).

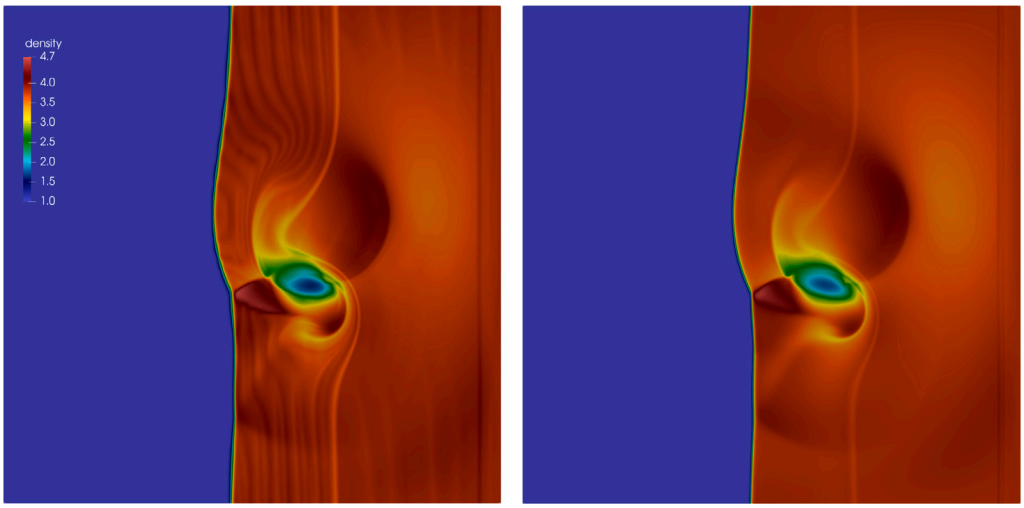

Fig. 15. Density contours of the (slow-moving) shock/vortex interaction problem.

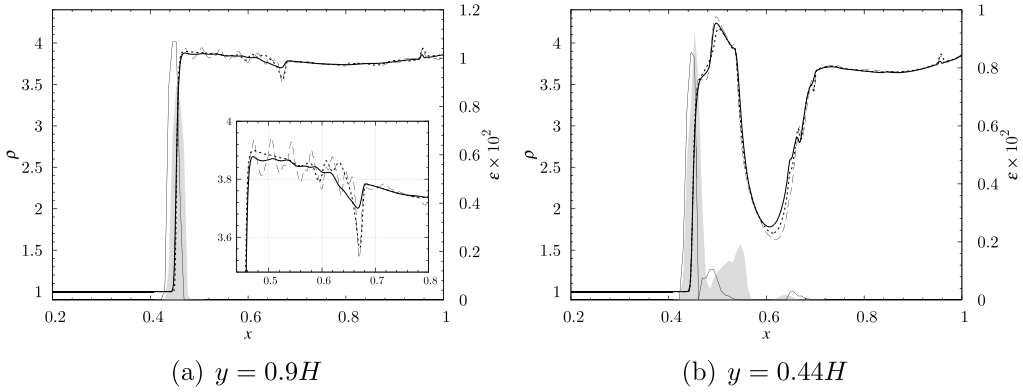

Fig. 16. Density and AV profiles across the shock: solid line, density with adaptive exponential smoothing, $K = 2$ (left axis); thin dashed line, density without exponential smoothing (left axis); dotted line, exponential smoothing at imposed shock speed $D = 5\%D^*$ (left axis); gray filled line, shock-capturing viscosity with adaptive exponential smoothing (right axis); thin solid line, shock-capturing viscosity without exponential smoothing (right axis).

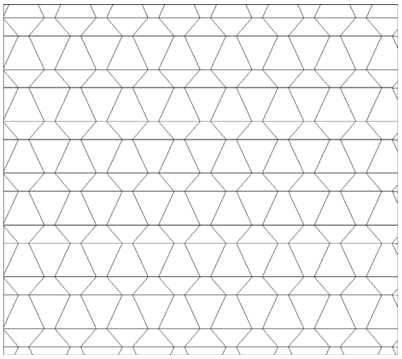

Fig. 17. Elements' pattern in the deformed mesh for the shock/vortex interaction problem.

where M_v represents the vortex Mach number whereas $\xi_v = (x - x_{v0})/R_v$ and $\eta_v = (y - y_{v0})/R_v$ are the spatial coordinates in the vortex reference frame normalized with respect to the vortex radius $R_v = 0.1H$ (x_{v0} and y_{v0} are the coordinates of the vortex center at the beginning of the simulation). The vortex Mach number M_v is set equal to 0.8.

The inviscid Euler equations are integrated in time with the RK45-SSP scheme with a 0.8 CFL coefficient. Shock-capturing is enforced by means of the Laplacian model, with the AV scaled on the maximum wave speed, with the parameter $C_\varepsilon = 1$. The positivity-preserving scheme is deactivated. In the case of the computation performed with the exponential smoothing, the number of centroids K of the adaptive power k-means calibration procedure is set equal to 2.

The density contours after the vortex has traversed the shock, at time $t = 0.08\,M_s a_L/H$ (a_L being the speed of sound upstream of the shock), are depicted in Fig. 15, whereas Fig. 16 shows the density profiles along two horizontal lines, one at $y = 0.9H$, in the region of most severe PSO and the other at $y = 0.44H$, passing through the vortex center. The results obtained with the adaptive exponential smoothing are compared against those obtained without exponential smoothing or with exponential smoothing, globally calibrated on the theoretical shock speed (this last case is only shown in Fig. 16).

As expected, due to the extremely low value of the shock speed, the computation performed without exponential smoothing is severely impacted by very strong PSO. These are almost completely suppressed when the exponential smoothing is active. Overall, the shock profile does not seem to be affected by the exponential smoothing in terms of its sharpness. Nonetheless, possibly due to the less localized distribution of the AV, the dissipation at and around the vortex appears to be slightly increased when the exponential smoothing is active.

Similar tests are performed using a distorted grid which features an arbitrary displacement of the grid points, as shown in Fig. 17. The resulting mesh is composed of highly-skewed, triangle-like elements,

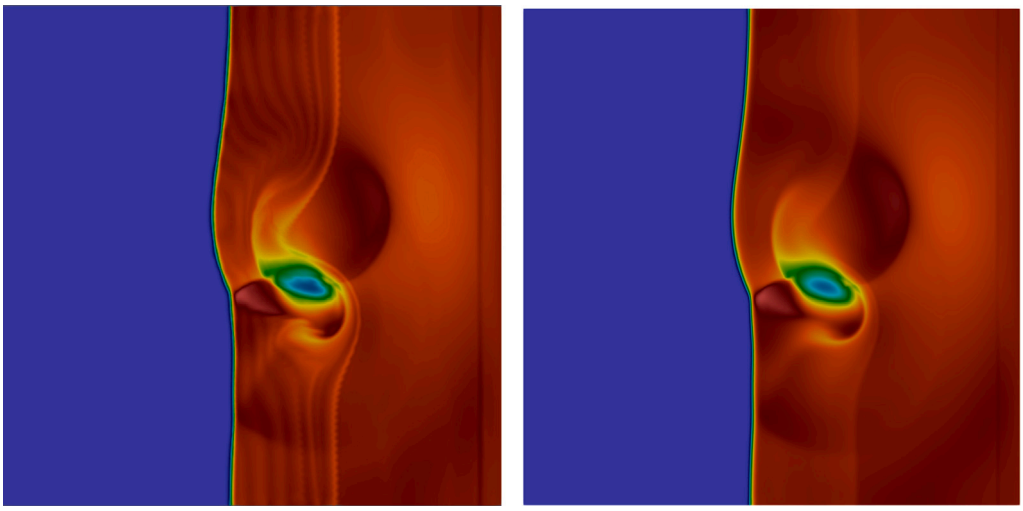

(a) no exponential smoothing (b) adaptive exponential smoothing

Fig. 18. Density contours of the (slow-moving) shock/vortex interaction problem on distorted meshes.

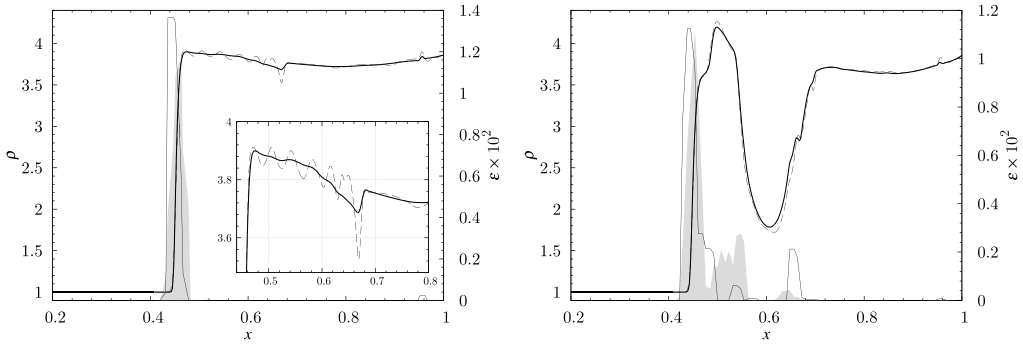

(a) $y = 0.9H$ (b) $y = 0.44H$

Fig. 19. Density and AV profiles across the shock for distorted meshes: solid line, density with adaptive exponential smoothing, $K = 2$ (left axis); thin dashed line, density without exponential smoothing (left axis); dotted line, exponential smoothing at imposed shock speed $\mathcal{D} = 5\%\mathcal{D}^*$ (left axis); gray filled line, shock-capturing viscosity with adaptive exponential smoothing (right axis); thin solid line, shock-capturing viscosity without exponential smoothing (right axis).

which were found to degrade vortex transport phenomena when a low order of accuracy was considered [34]. This type of grid is of interest to assess the accuracy of the method on meshes more representative of applications. The resolution in terms of number of elements and order of accuracy is kept the same as the cartesian computations. Fig. 18 displays the density contours at the final time both with and without exponential smoothing for the AV formulation. It is observed that the computation without exponential smoothing displays strong wiggles of numerical nature past the shock. The exponential smoothing is effectively smearing most of the oscillating features, yielding a smooth density field which confirms the performance of the new method on distorted grids. In order to provide a more quantitative view of the improvement, the density profiles are extracted at the locations $x = 0.9H$ and $x = 0.44H$, see Fig. 19. As was observed on Cartesian meshes, the exponential smoothing clearly removes the high amplitude PSO that the computation without exponential smoothing displays. It is also interesting to notice that the density contours, as well as the density and AV profiles, are almost identical between the Cartesian and distorted meshes computations with exponential smoothing, emphasizing the relative insensitivity of the proposed methodology to grid distortion.

6.5. Double mach reflection problem

The double Mach reflection (DMR) of a strong shock problem, originally proposed by Woodward and Colella [35], is here used to assess the proposed methodology on a configuration in which, *a priori*, there is no major problem of PSO. It involves a Mach 10 shock, propagating in gas at rest, colliding with a ramp inclined at 30° with respect to the shock propagation direction. Upon collision, a self-similar structure is formed over the ramp, which consists of two triple points, each one comprising three shocks and a slip line [36, Fig. 3.9, p. 149].

Notice that the shock propagation velocity is indeed quite high and significant PSO are not expected to become an issue. Hence, for this configuration, the exponential smoothing approach would not be strictly required. Nonetheless, it is useful to check the impact of the exponential smoothing on the development of the self-similar structure that characterizes the flow setup.

The typical setup for the DMR problem consists of a rectangular computational domain, which is aligned with the ramp and in which an oblique Mach 10 shock is initialized at 60° with respect to the horizontal axis. As a result of the numerical setup, several numerical artifacts can impact the solution [35,37], namely: (a) numerical artifacts due to the initialization of an oblique shock over a non-aligned Cartesian grid, (b) the "starting error" due to the shock being initialized as a sharp profile and (c) a numerical artifact originating on the top boundary due to the objective difficulty in prescribing a *clean* boundary condition for an oblique shock propagating along a horizontal boundary line.

To overcome the above difficulties and taking advantage of the unstructured nature of the solver used in the present study, the computational mesh adopted here is a parallelogram as sketched in Fig. 20.

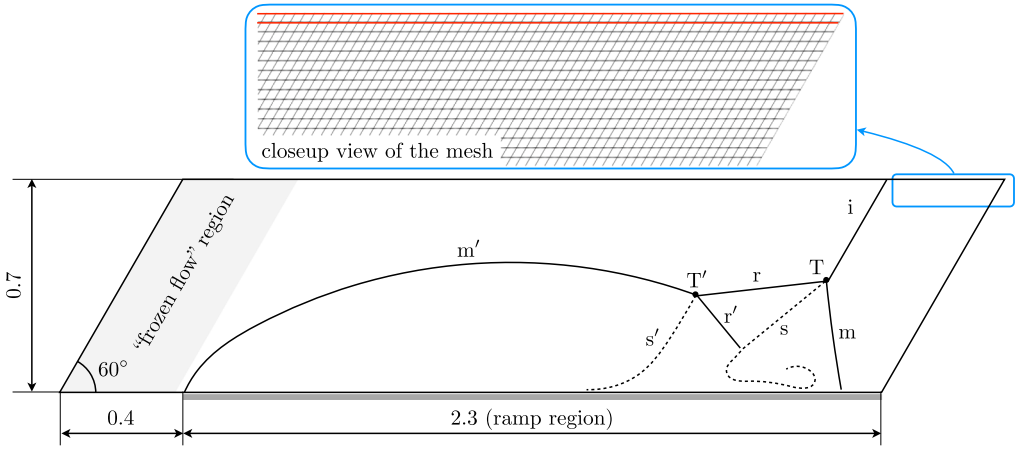

Fig. 20. Sketch of DMR problem computational domain and solution (red lines in the closeup view indicate periodic interfaces): i, incident shock; T, first triple point; m, primary Mach stem; r, primary reflected shock; s, primary slip line; T', second triple point; m', secondary Mach stem; r', secondary reflected shock; s', secondary slip line.

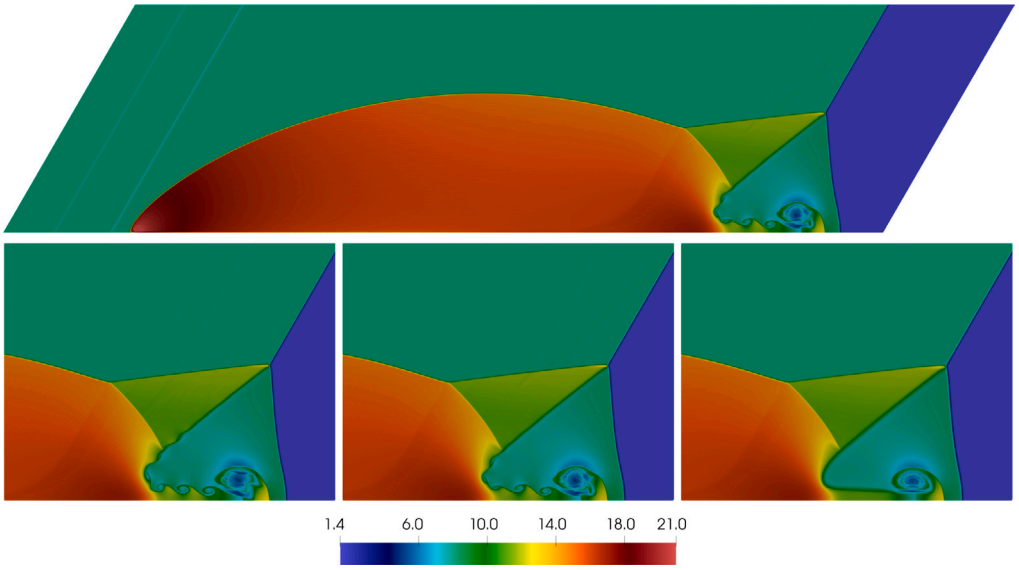

Fig. 21. Density contours of DMR problem: bottom left, no exponential smoothing; top and bottom middle, $C_{ea} = 3.0$; bottom right, $C_{ea} = 1.0$.

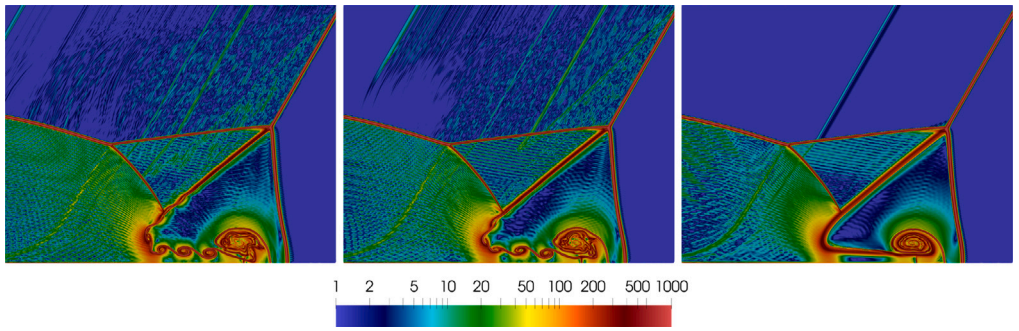

Fig. 22. Norm of density gradient contours of DMR problem: left, no exponential smoothing; middle, $C_{ea} = 3.0$; right, $C_{ea} = 1.0$.

The main shock (i), initially located at the left boundary, is hence perfectly aligned with the mesh. To deal with the spurious waves, which are formed as the scheme resolves the initial sharp profile and which trail behind the smoothed shock, by the time the shock reaches the leading edge of the ramp, these waves are isolated from the flow by setting to zero the residuals within the “frozen flow” region as indicated in Fig. 20. This way, the perturbations remain *stuck* upstream of the ramp and cannot interfere anymore with the DMR structure. Concerning the numerical artifact from the top boundary, a similar approach as the one proposed by Vevek, et al. [37] is adopted. Accordingly, the top boundary is made periodic with respect to a line just below it. In practice, taking advantage of the DFE structure of the solver and the deformation imposed to the mesh, this artificial periodicity translates, straightforwardly, in making the boundary interface of every boundary

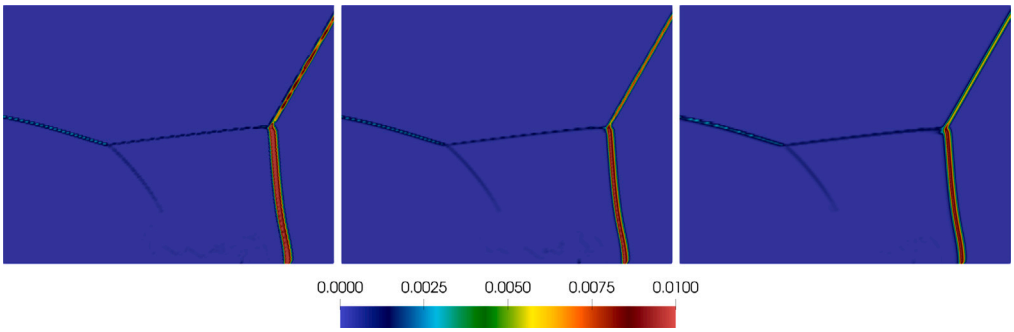

Fig. 23. Artificial viscosity contours of DMR problem: left, no exponential smoothing; middle, $C_{ea} = 3.0$; right, $C_{ea} = 1.0$.

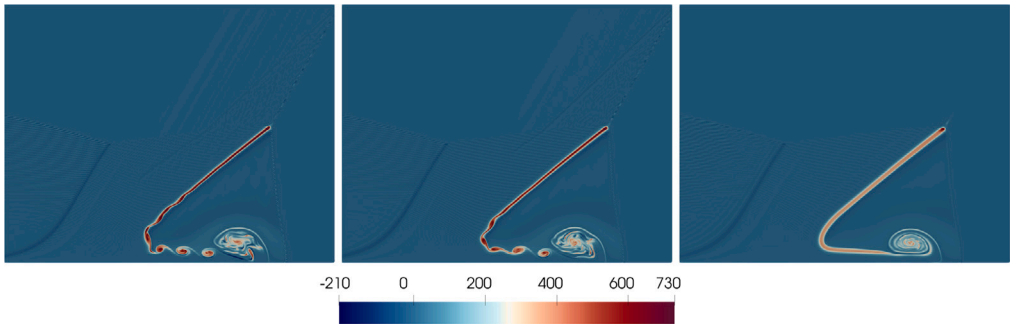

Fig. 24. Vorticity contours of DMR problem: left, no exponential smoothing; middle, $C_{ea} = 3.0$; right, $C_{ea} = 1.0$.

element periodic with respect to the opposing interface within the fluid (*cf.* Fig. 20). Notice that the same approach is initially used at the bottom boundary of the frozen flow region before switching to an imposed post-shock state when the shock reaches the ramp leading edge. This switch of boundary type is needed to keep the foot of the second Mach stem (m') attached to the ramp leading edge. Finally, the ramp boundary condition is set as a free-slip wall and the right boundary is set as an outflow.

Fifth-order computations have been performed with and without adaptive exponential smoothing ($K = 2$) on 810×210 elements (4.2 M DoF) and a mesh resolution $h = 1/300$. The simulation was started with the shock at the left boundary and stopped at a time $t = 0.2$ after the moment the shock reaches the leading edge of the ramp. It is worth stressing again that this test features an extremely fast shock. Such a fast shock does not promote PSO and, therefore, the exponential smoothing is not strictly required. To check nonetheless the impact of the exponential smoothing on the solution, two different values of the smoothing parameter C_{ea} have been tested: a value of 1.0, which promotes a relatively mild smoothing with a non-negligible penalty in AV localization, and a higher value of 3.0, which produces marginal smoothing and AV delocalization (*cf.* Section 4.1). The AV has been injected via the Laplacian model. It is worth noting that, possibly due to the topology of the mesh, mild carbuncle phenomena had been observed close to the first triple point (T). To prevent the onset of these perturbations, the AUFS numerical flux has been adopted [38].

The results are reported in Figs. 21–24, showing, respectively, the contours of density, density gradient norm, artificial viscosity and vorticity. As anticipated, no PSO are visible in the density contours, even without exponential smoothing. All the expected features of the self-similar DMR structure are well reproduced in all the computations. In particular, the secondary slip line (s'), one of the main challenges for numerical schemes in this setup, although very subtle, is clearly visible in the density map.

A noticeable difference in the results is in the destabilization of the primary slip line (s): when the exponential smoothing is active, the instability is slightly delayed ($C_{ea} = 3.0$) or completely suppressed ($C_{ea} = 1.0$). The reason behind this behavior appears to be twofold.

First of all, the reduction ($C_{ea} = 3.0$) or the complete suppression ($C_{ea} = 1.0$) of the spurious numerical perturbations which are clearly visible in Fig. 22 behind the leading shock, and possibly responsible for promoting the destabilization of the slip line. By inspection of Fig. 23, in particular, the origin of these perturbations can be clearly identified in the isolated peaks of AV which are visible along the leading shock. Such peaks of AV have been already observed in one-dimensional problems (*cf.* Fig. 1b) and, as an evident source of (unphysical) destabilization, in a two-dimensional Riemann problem (*cf.* Fig. 13). These *spots* of AV are mitigated or completely suppressed depending on the strength of the applied exponential smoothing, as it can be seen in the middle and right pictures of Fig. 23.

Secondly, the increase in the width of the slip line and the consequent reduction of the relevant vorticity (see Fig. 24) which, in turn, make the slip line less sensitive to external perturbations. Such increase in the width of the primary slip line is particularly evident in Fig. 22 for the case with $C_{ea} = 1.0$. As a result, compared to the baseline case without exponential smoothing, if the simulation with $C_{ea} = 3.0$ produced a reduction of about 4%–12% in the vorticity on the primary slip line, in the computation with $C_{ea} = 1.0$, this reduction was more significant and around 24%–48%.

The culprit of the thickening of the primary slip line, a phenomenon that was already observed in the two-dimensional Riemann problems in Section 6.3, can be identified in the contours of AV in Fig. 23. In fact, while both simulations have negligible values of AV in the neighborhood of the slip line and localize AV where the shocks are present, a slight excess of viscosity can be seen, for the case with $C_{ea} = 1.0$, right below the first triple point, where the primary slip line originates. Such increased viscosity is most probably responsible for the thickening of the slip line during the evolution of the self-similar DMR structure, especially during the early stages of its development, when, due to the small size of the DMR structure, the mesh resolution is comparatively extremely low.

Ultimately, whether or not a slip line within an inviscid fluid is supposed to destabilize, spontaneously, in the absence of an external

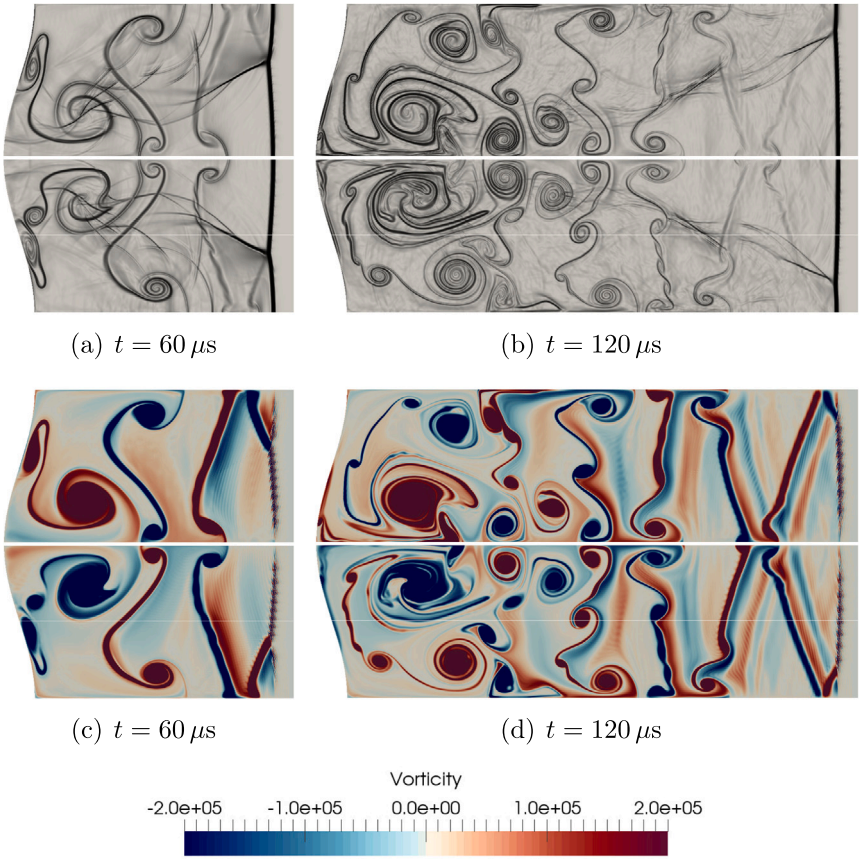

Fig. 25. Numerical Schlieren and vorticity contours of shock/wavy-wall interaction problem with (top half of the plots) and without (bottom half of the plots) exponential smoothing. The thin line indicate the location along which the profiles in Fig. 26 are extracted.

Table 7
Mesh parameters and flow setup. The velocities are indicated in the laboratory reference frame, oriented from the left to the right.

Domain size	$L \times H$	10.0×2.0	cm^2
Number of elements	$N_x \times N_y$	600×140	–
Discretization order	n	5	–
Degrees of freedom	DoF	2.1×10^6	–
Element resolution	Δ/n	28.6–33.3	μm
Wavy wall amplitude	A_{ww}	1.0	mm
Wavy wall wavelength	λ_{ww}	2.0	cm
Specific heat ratio	γ	1.15	–
Viscosity (at 291.15 K)	μ_0	1.827×10^{-5}	kg/(ms)
Initial left state	ρ_1	1.208	kg/m^3
	p_1	1.0	atm
	u_1	0.0	m/s
Initial right state	ρ_2	11.29	kg/m^3
	p_2	26.68	atm
	u_2	−1386.88	m/s
Incident shock Mach	M_i	5.0	–
Incident shock speed	D_i	−1553.0	m/s
Reflected shock Mach	M_r	3.16	–
Reflected shock speed	D_r	271.5	m/s

perturbation remains however a legitimate question. Indeed, if the destabilization of the primary slip line is commonly regarded as an evidence of low dissipation and good resolving power of the numerical scheme, it is also true that the vast majority of the results reported in the literature for the DMR problem have been obtained in numerical setups which are definitely not free from spurious numerical perturbations (*e.g.*, Cartesian grids not aligned with the shock, starting errors, boundary perturbations).

6.6. Shock/wavy-wall interaction in the Newtonian limit

The last test case concerns the reflection of a strong shock in the Newtonian limit by a smoothly curved wall. This test is identical to the one recently reported in [39]. The results obtained using the proposed adaptive exponential smoothing are compared with those reported in [7,8], where PSO reductions were also attained using characteristic based sensors. It is worth noting that the same characteristic (density/acoustic) sensor used in the above mentioned works is also used in the present tests. Hence, the comparison aims at assessing the impact of the exponential smoothing on this particular flow configuration.

The physical parameters of the simulation are summarized in Table 7 for convenience. The Navier–Stokes equations are integrated using a 5th-order SD discretization. Time integration is performed with the RK45-SSP, with 0.8 CFL and Fourier coefficients. To avoid instabilities during the establishment of the initial shock front, the positivity-preserving scheme [26] is active. The left no-slip wall has a sinusoidal shape with 1 mm amplitude and 2 cm wavelength. The right boundary is set as an inflow, whose parameters correspond to the compressed (shocked gases) state of the incident shock. The top and bottom boundaries are periodic. Shock-capturing is enforced by means of a Laplacian term, whose viscosity is scaled on the maximum speed of sound. The scaling parameter C_ε is set at 0.75, which was found to be an optimal choice for this test [7].

Numerical Schlieren[7] images and vorticity contours obtained 60 μs and 120 μs after the leading shock impacts the wavy-wall are depicted

[7] As per Ref. [40], numerical Schlieren images are obtained as the contours of

$$S = 0.8 \exp\left(\frac{14.0\,\|\nabla\rho\|}{3 \times 10^5}\right).$$

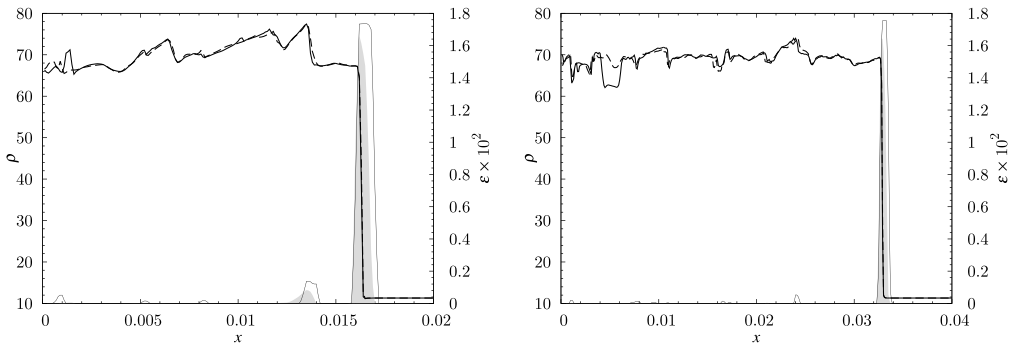

Fig. 26. Density and AV profiles along the line $y = -0.5$ cm: solid line, adaptive exponential smoothing, $K = 2$ (left axis); dashed line, no exponential smoothing (left axis); gray filled line, shock-capturing viscosity with adaptive exponential smoothing (right axis); thin solid line, shock-capturing viscosity without exponential smoothing (right axis).

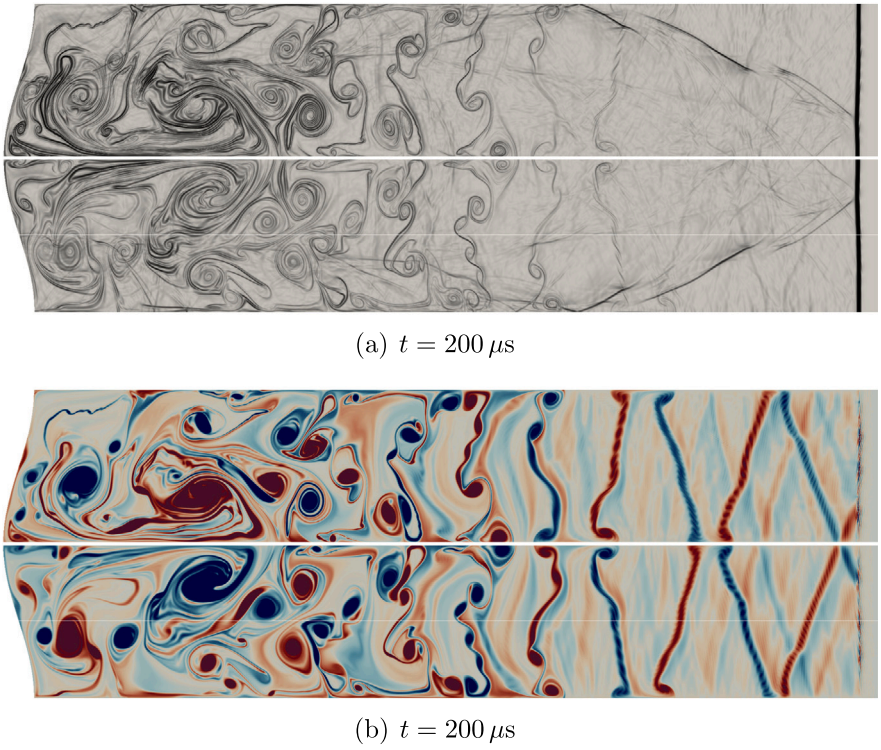

Fig. 27. Numerical Schlieren and vorticity contours of shock/wavy-wall interaction problem with (top half of the plots) and without (bottom half of the plots) exponential smoothing (vorticity scale as per Fig. 25).

in Fig. 25. In each snapshot, the top half of the contour plot refers to the simulation performed with the adaptive exponential smoothing with the power k-means procedure at $K = 2$. Notice that the results obtained with $K = 3$ (not shown) are indistinguishable from those reported here; this further confirms the robustness of the power k-means approach in respect of the selected value of K. The bottom half refers to the baseline computation from [7]. For the same simulated times, density and AV profiles along the line $y = -0.5$ cm (indicated in the contour plots with a thin white line in the bottom half of each snapshot) are also compared in Fig. 26.

As it can be seen, the results with and without exponential smoothing are marginally different in this case. Indeed, the characteristic sensor promotes already PSO reductions and the exponential smoothing cannot do much better. When the exponential smoothing is active, some of the slip lines—or contact discontinuities—characterizing the shocked gas region (see [25] for details) do not destabilize as observed in the baseline computation. Although a slight reduction of AV localization at the shock is expected when using exponential smoothing, this phenomenon does not appear to be due to excessive residual dissipation. Indeed, comparing the AV contours of the two computations (not shown), the exponential smoothing promotes, overall, a decrease of the level of residual dissipation within the shocked gas and a significant reduction of isolated peaks of AV. Analogous to what was observed for the two-dimensional Riemann problem in Section 6.3, it seems then plausible to correlate the observed behavior to the suppression of localized and spurious destabilization mechanisms by the AV.

Finally, the numerical Schlieren and vorticity contours for $t = 200$ μs are depicted in Fig. 27 where previous conclusions for this particular test are confirmed: (a) the exponential smoothing has a negligible impact in the region just behind the leading shock, where transverse shocks and slip lines are originated; ((b)) far from the shock, the overall level of AV and the occurrence of localized and isolated AV peaks are reduced thanks to the time delay intrinsic to the exponential smoothing approach; (c) the different and, possibly, better behavior of the AV

away from the leading shock, has an impact on the long term evolution of the flow.

7. Conclusion

A novel approach to mitigate, or even suppress, post-shock oscillations has been proposed. The method is here formalized within the framework of discontinuous finite element methods using shock-capturing techniques with modal sensors. These methods—and possibly other methods relying on shock-capturing methodologies—are affected by post-shock oscillations, which can become extremely severe in the case of slow-moving shocks. The main culprit of these oscillations has been identified in the inherent inhomogeneities in the modal detector (differences in the way the same discontinuity is detected depending upon its position within the elements) which, in turn, promote unsteadiness of the injected artificial viscosity and its sub-optimal distribution around the shocks.

In order to overcome and mitigate the problem, rather than trying to improve the behavior of the detector and the spatial distribution of the artificial viscosity in the neighborhood of discontinuities, as already attempted with some success by the first author using characteristic sensors [7,8], a completely different approach has been proposed. This aims at improving the time behavior of the artificial viscosity via a suitably calibrated delay, which is introduced to prevent the artificial viscosity from undergoing excessive changes during the time a discontinuity crosses an element. This form of time regularization is achieved by means of the exponential averaging technique, where the relevant smoothing factor is computed from the element's geometry and the propagation velocity of the discontinuity.

A methodology has also been proposed to estimate the propagation velocity of one or even several shocks in the domain, which can then be used to properly tune the exponential averaging approach on every shock. To this end, to tackle the rather challenging task of computing the propagation velocity of (quite well possibly) badly behaved discontinuous and oscillatory signals (the shocks), the power k-means clustering technique—commonly used in machine learning—is adopted. This method enables the identification of each shock in the flow and provides the best estimate of the relevant propagation velocities.

The proposed technique depends, essentially, on a single main tuning parameter (C_{ea}), which sets the time delay of the shock sensor response. For the majority of flows, with mild-to-medium post-shock oscillations, this tuning parameter can be safely set to a unitary value, a setting which is expected to already produce significant reductions (or even the suppression) of all spurious oscillations. For more challenging configurations, with very slow shocks and severe post-shock oscillations, lower values of the tuning parameter (down to a minimum recommended value of 0.5) might be advisable. Flows involving stationary or fast-moving shocks, for which post-shock oscillations do not represent an issue, can be safely computed with higher values of the tuning parameter (*i.e.*, $C_{\text{ea}} \sim 3.0$) or switching off the proposed exponential averaging technique.

Several results on one- and two-dimensional tests have demonstrated that the proposed strategy achieves, quite consistently, reductions in post-shock oscillations from about 70%, up to 100% (complete suppression), without noticeable reductions in shock sharpness. On this regard, in particular, due to the potential reduction in artificial viscosity that the exponential averaging may cause, some tests have even shown sharper (and post-oscillations free) shocks. On the other hand, a penalty has been observed in general in terms of spatial localization of the artificial viscosity which, in some cases, can promote localized increased dissipation. This phenomenon, which is intimately connected to the introduction of a time delay by the exponential averaging approach, is particularly evident when using low orders and low resolutions or, in some cases, in two-dimensional tests, where slip lines issuing from triple points result somewhat thickened, and can be mitigated, if deemed necessary, by making the exponential averaging less pronounced (the only caveat being, of course, to give in some reduction in post-shock oscillations).

On this regard, although it might be tempting to conclude that post-shock oscillations reductions are essentially related to an increase of the injected artificial dissipation, it is important to remember that, when the exponential smoothing is not used, a simple increase of the level of artificial viscosity will rather lead to an amplification of post-shock oscillations [7,8]. This is because the modal sensor itself is causing post-shock oscillations and a higher setting for the artificial viscosity will determine even stronger perturbations in the flow. As a matter of fact, no post-shock oscillations reductions are possible just by increasing how much artificial viscosity is injected without taking very good care of how this is injected. In situations where post-shock oscillations represent a major concern, despite the above-mentioned penalty in spatial localization, the proposed technique is, to the authors' knowledge, the only effective way to significantly reduce (or completely suppress) them.

Future developments and improvements of the proposed methodology shall focus on reliable techniques to automatically deactivate the exponential smoothing—or better adapt the relevant tuning parameter—in flows involving fast-moving shocks accompanied by negligible post-shock oscillations.

Finally, the proposed methodology can be applied to the whole class of high-order discontinuous finite elements methods (possibly, in other methods too, such as finite volumes) to significantly suppress post-shock oscillations on an arbitrary number of different shocks in the flow.

CRediT authorship contribution statement

Guido Lodato: Conceptualization, Methodology, Software, Validation, Investigation, Writing – original draft, Writing – review & editing. **Luc Vervisch:** Conceptualization, Methodology, Writing – review & editing. **Jean-Baptiste Chapelier:** Conceptualization, Methodology, Validation, Investigation, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The first author would like to thank Catherine Vesco for fruitful discussions about the results and the structure of the manuscript. The use of the SD solver originally developed by Antony Jameson's group at Stanford University is gratefully acknowledged. This work was granted access to the high-performance computing resources of CRIANN.

Funding

Financial support to the first and second authors was provided by the French National Research Agency (ANR) under grant number ANR-18-CE05-0030.

References

[1] Cook AW, Cabot WH. A high-wavenumber viscosity for high-resolution numerical methods. J Comput Phys 2004;195(2):594–601.
[2] Persson P-O, Peraire J. Sub-cell shock capturing for discontinuous Galerkin methods. In: AIAA P. 2006–112. 2006, p. 1–13, 44th AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV, Jan. 9–12, 2006.
[3] Persson P-O. Shock capturing for high-order discontinuous Galerkin simulation of transient flow problems. In: AIAA P. 2013–3061. 2013, p. 1–9, 21st AIAA Computational Fluid Dynamics Conference, San Diego, CA, Jun. 24–27, 2013.

[4] Klöckner A, Warburton T, Hesthaven J. Viscous shock capturing in a time-explicit discontinuous Galerkin method. Math Model Nat Phenom 2011;6(3):57–83.

[5] Tonicello N, Lodato G, Vervisch L. Entropy preserving low dissipative shock capturing with wave-characteristic based sensor for high-order methods. Comput Fluids 2020;197:104357. http://dx.doi.org/10.1016/j.compfluid.2019.104357.

[6] Arora M, Roe PL. On postshock oscillations due to shock capturing schemes in unsteady flows. J Comput Phys 1997;130(1):25–40.

[7] Lodato G. Characteristic modal shock detection for discontinuous finite element methods. Comput Fluids 2019;179:309–33. http://dx.doi.org/10.1016/j.compfluid.2018.11.008.

[8] Lodato G. Corrigendum to the article "Characteristic modal shock detection for discontinuous finite element methods". Comput Fluids 2019;193:104245. http://dx.doi.org/10.1016/j.compfluid.2019.104245.

[9] Kopriva D, Kolias J. A conservative staggered-grid Chebyshev multidomain method for compressible flows. J Comput Phys 1996;125(1):244–61.

[10] Sun Y, Wang Z, Liu Y. High-order multidomain spectral difference method for the Navier-Stokes equations on unstructured hexahedral grids. Commun Comput Phys 2007;2(2):310–33.

[11] Jameson A. A proof of the stability of the spectral difference method for all orders of accuracy. J Sci Comput 2010;45(1):348–58.

[12] Cockburn B, Shu C. The local discontinuous Galerkin finite element method for convection-diffusion systems. SIAM J Numer Anal 1998;35:2440–63.

[13] Cockburn B, Shu C. The Runge-Kutta discontinuous Galerkin finite element method for conservation laws V: Multidimensional systems. J Comput Phys 1998;141:199–224.

[14] Hesthaven JS, Warburton T. Nodal discontinuous Galerkin methods: Algorithms, analysis, and applications. Springer Science+Business Media, LLC; 2008.

[15] Wang ZJ. Spectral (finite) volume method for conservation laws on unstructured grids: Basic formulation. J Comput Phys 2002;178(1):210–51.

[16] Liu Y, Vinokur M, Wang Z. Spectral (finite) volume method for conservation laws on unstructured grids V: extension to three-dimensional systems. J Comput Phys 2006;212(2):454–72.

[17] Huynh H. A flux reconstruction approach to high-order schemes including discontinuous Galerkin methods. In: AIAA P. 2007–4079. 2007, p. 1–42, 18th AIAA Computational Fluid Dynamics Conference, Miami, FL, Jun. 25–28, 2007.

[18] Vincent P, Castonguay P, Jameson A. A new class of high-order energy stable flux reconstruction schemes. J Sci Comput 2010;47(1):1–23.

[19] Liu Y, Vinokur M, Wang Z. Spectral difference method for unstructured grids I: basic formulation. J Comput Phys 2006;216(2):780–801.

[20] Liang C, Jameson A, Wang Z. Spectral difference method for compressible flow on unstructured grids with mixed elements. J Comput Phys 2009;228(8):2847–58.

[21] Roe P. Approximate Riemann solvers, parameter vectors, and difference schemes. J Comput Phys 1981;43:357–72.

[22] Harten A. High resolution schemes for hyperbolic conservation laws. J Comput Phys 1983;49(3):357–93.

[23] Spiteri R, Ruuth S. A new class of optimal high-order strong-stability-preserving time discretization methods. SIAM J Numer Anal 2002;40(2):469–91.

[24] Barter GE, Darmofal DL. Shock capturing with PDE-based artificial viscosity for DGFEM: Part I. Formulation. J Comput Phys 2010;229(5):1810–27.

[25] Lodato G, Vervisch L, Clavin P. Direct numerical simulation of shock wavy-wall interaction: analysis of cellular shock structures and flow patterns. J Fluid Mech 2016;789:221–58. http://dx.doi.org/10.1017/jfm.2015.731.

[26] Zhang X, Shu C-W. On positivity-preserving high order discontinuous Galerkin schemes for compressible Euler equations on rectangular meshes. J Comput Phys 2010;229(23):8918–34.

[27] Sod GA. A survey of several finite difference methods for systems of nonlinear hyperbolic conservation laws. J Comput Phys 1978;27(1):1–31.

[28] Xu J, Lange K. Power k-means clustering. In: International conference on machine learning. PMLR; 2019, p. 6921–31.

[29] Zhang B, Hsu M, Dayal U. K-harmonic means - a data clustering algorithm. Technical Report HPL-1999-124, Palo Alto, CA: Hewlett-Packard Laboratory; 1999.

[30] Shu C-W, Osher S. Efficient implementation of essentially non-oscillatory shock-capturing schemes, II. J Comput Phys 1989;83(1):32–78.

[31] Lax PD, Liu X-D. Solution of two-dimensional Riemann problems of gas dynamics by positive schemes. SIAM J Sci Comput 1998;19(2):319–40.

[32] Liska R, Wendroff B. Comparison of several difference schemes on 1D and 2D test problems for the Euler equations. SIAM J Sci Comput 2003;25(3):995–1017.

[33] Taylor G. On the dissipation of eddies. In: The scientific papers of Sir Geoffrey Ingram Taylor, Vol. 2. Cambridge University Press; 1918, p. 96–101.

[34] Chapelier J-B, Lodato G. Study of the spectral difference numerical dissipation for turbulent flows using unstructured grids. Flow Turbul Combust 2017;99(3):643–64. http://dx.doi.org/10.1007/s10494-017-9847-5.

[35] Woodward P, Colella P. The numerical simulation of two-dimensional fluid flow with strong shocks. J Comput Phys 1984;54(1):115–73.

[36] Ben-Dor G. Shock wave reflection phenomena. Springer; 2007.

[37] Vevek U, Zang B, New T. On alternative setups of the double Mach reflection problem. J Sci Comput 2019;78(2):1291–303.

[38] Sun M, Takayama K. An artificially upstream flux vector splitting scheme for the Euler equations. J Comput Phys 2003;189(1):305–29.

[39] Lodato G, Vervisch L, Clavin P. Numerical study of smoothly perturbed shocks in the Newtonian limit. Flow Turbul Combust 2017;99(3):887–908. http://dx.doi.org/10.1007/s10494-017-9830-1.

[40] Hadjadj A, Kudryavtsev A. Computation and flow visualization in high-speed aerodynamics. J Turbul 2005;6(16):1–25.