INTERNATIONAL JOURNAL FOR NUMERICAL METHODS IN FLUIDS
Int. J. Numer. Meth. Fluids 2016; **82**:398–416
Published online 2 March 2016 in Wiley Online Library (wileyonlinelibrary.com). DOI: 10.1002/fld.4223

Dilation-based shock capturing for high-order methods

David Moro*,†, Ngoc Cuong Nguyen and Jaime Peraire

Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, 77 Mass Ave, Cambridge, MA 02139, USA

SUMMARY

In this paper, we introduce a shock-capturing artificial viscosity technique for high-order unstructured mesh methods. This artificial viscosity model is based on a non-dimensional form of the divergence of the velocity. The technique is an extension and improvement of the dilation-based artificial viscosity methods introduced in Premasuthan *et al.* [15] and further extended in Nguyen and Peraire [27]. The approach presented has a number attractive properties including non-dimensional analytical form, sub-cell resolution, and robustness for complex shock flows on anisotropic meshes. We present extensive numerical results to demonstrate the performance of the proposed approach. Copyright © 2016 John Wiley & Sons, Ltd.

Received 13 September 2015; Revised 15 January 2016; Accepted 24 January 2016

KEY WORDS: compressible flow; discontinuous Galerkin; supersonic; transonic; implicit; nonlinear solvers

1. INTRODUCTION

The development of robust, accurate, and efficient methods for the numerical solution of hyperbolic nonlinear systems of conservation laws in complex geometries is a topic of considerable importance. Hyperbolic partial different equations (PDEs) govern a wide range of physical phenomena and arise in several areas of applied mathematics and mechanics such as fluid dynamics, magnetohydrodynamics, multiphase flow, population dynamics, and traffic flow. A distinguishing feature of the solutions of nonlinear hyperbolic systems is the potential for discontinuities, or shock waves, to develop and propagate even if initial and boundary data are smooth. The presence of such discontinuities is a major challenge when attempting to compute physically correct and stable solutions of hyperbolic conservation laws. Although significant progress has been made over the past decades, the numerical approximation of hyperbolic PDEs remains an active research area.

The compressible Euler and Navier–Stokes equations are one of the most important and widely investigated systems of hyperbolic conservation laws. High-order methods have gained increasing attention in recent years for the solution of the compressible Euler and Navier–Stokes equations because of the need for high-order accuracy in applications such as direct numerical simulation, large eddy simulation, and computational aeroacoustics. For these applications, current research activity is aimed at the development of efficient and robust high-order methods, and their application to real-world problems. This includes the numerical treatment of shocks, which is the focus of this paper.

One of the most common approaches to deal with shocks in the context of high-order methods is the use of flux/slope limiting techniques. The early Essentially non-oscillatory (ENO) and Weighted ENO (WENO) schemes [1–3] for finite volume discretizations choose the reconstruction stencil to

*Correspondence to: David Moro, Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, 77 Mass Ave, Cambridge, MA 02139, USA.
†E-mail: dmoro@mit.edu

minimize oscillations around discontinuities. For discontinuous Galerkin (DG) methods, Cockburn and Shu [4, 5] developed the Runge–Kutta discontinuous Galerkin scheme, which limits the high-order representation of the solution inside each element. In this same fashion, Krivodonova [6] proposed to detect shocks based on jumps at the interfaces between elements and later presented a limiting strategy based on monitoring the high-order coefficients of the solution [7]. A major drawback of most limiters is that they are non-smooth and thus present difficulties for implicit time-stepping methods.

A different approach to shock capturing, dating back to the work of Von Neumann and Richtmeyer [8] in the 1950s, is the use of artificial viscosity. Artificial viscosity has been widely used in finite volume methods [9, 10], streamline upwind Petrov–Galerkin methods [11], and spectral methods [12]. For higher order finite difference and finite volume methods, Cook and Cabot [13, 14] incorporated artificial viscosity by scaling the physical viscosity terms such as the dynamic viscosity, bulk viscosity, and thermal conductivity. The magnitude of the scaling factors is determined from the strain rate tensor and the internal energy. This approach was later adopted by other authors in the context of compressible turbulence simulations, and by Premasuthan *et al.* [15, 16] for spectral difference method using only the dilation part of the shock sensor proposed by Bhagatwala and Lele [17].

Artificial viscosity has also been employed with DG methods to capture shocks by applying artificial viscosity to all the equations and using artificial viscosity coefficients, or sensors, based on element residuals [18–22]. However, these residual-based methods tend to present some robustness issues when used with high-order methods and strong shocks.

Persson and Peraire [23] introduced a sub-cell shock-capturing artificial viscosity approach for high-order DG methods where the artificial viscosity coefficient is based on the smoothness of the computed density, which in turn is determined by the rate of decay of the coefficients of an orthogonal expansion of the density field. Later, Klockner *et al.* [24] and Persson [25] improved this model by considering a continuous piecewise linear reconstruction of the artificial viscosity coefficient. Barter and Darmofal [26] proposed a PDE-based artificial viscosity model appended to the system of governing equations to obtain a smoother artificial viscosity coefficient at the expense of solving an additional PDE.

In this paper, we present a shock-capturing artificial viscosity technique for high-order methods. The artificial viscosity coefficient is based on a non-dimensional form of the divergence of the velocity and is designed to obtain sub-cell shock resolution with high-order discretizations. The technique is an extension and improvement of the dilation-based artificial viscosity methods considered by Nguyen and Peraire [27] and Premasuthan *et al.* [15, 16]. The present technique has a smooth non-dimensional shock sensor with analytical form and is robust for complex flows on highly anisotropic meshes. Moreover, the technique can be implemented in any high-order numerical scheme in a straightforward manner. In this paper, we employ the hybridizable discontinuous Galerkin (HDG) method [28–31] with the proposed artificial viscosity model for the Euler equations. Compared with other DG methods, the HDG method possesses some crucial advantages in terms of global degrees of freedom count and convergence properties, as documented in recent work [30–33].

The paper is organized as follows. In Section 2, we introduce the governing equations, present the numerical method used, and describe the proposed artificial viscosity technique. In Section 3, we present extensive numerical results in 1D to calibrate and verify the model. In Section 4, we consider the multi-dimensional problem and show a variety of 2D flows. Finally, in Section 5, we present conclusions and discuss some directions for future research.

2. GOVERNING EQUATIONS AND DISCRETIZATION

We are interested in the solution of the equations that govern compressible flows. In the inviscid case, these are given by the Euler equations, which, written in conservative form, read as follows:

$$\frac{\partial \mathbf{u}}{\partial t} + \nabla \cdot \mathbf{F}(\boldsymbol{u}) = 0, \tag{1}$$

DOI: 10.1002/fld

where

$$\mathbf{u} = \begin{bmatrix} \rho \\ \rho v_i \\ \rho E \end{bmatrix}, \qquad \mathbf{F}(\boldsymbol{u}) = \begin{bmatrix} \rho v_i \\ \rho v_i v_j + P\delta_{ij} \\ \rho v_i H \end{bmatrix}. \tag{2}$$

Here, ρ represents the density, v_i is the component of the velocity in the i-th cartesian space dimension, E is the total specific energy, $H = E + P/\rho$ is the total specific enthalpy, and t is time. The pressure, P, the density, and the temperature, T, obey the ideal gas law:

$$P = \rho R T, \tag{3}$$

where R is the specific gas constant for air.

The governing equations in conservative form are discretized using a high-order HDG method [27–31]. This method was developed to reduce the number of globally coupled unknowns in DG discretizations and produce an approximation to the gradient of the solution that converges optimally for diffusion problems. In addition, the HDG method retains certain desirable features of DG methods such as arbitrarily high-order approximation, stability for convective operators, and the ability to deal with unstructured meshes. In this work, we follow closely the HDG discretization of the Euler equations introduced by Peraire *et al.* [30].

For the steady-state simulations presented in this paper, we use time continuation with a backward Euler scheme of *global* time step Δt. This requires the solution of a nonlinear system of algebraic equations at each time step, which is carried out using Newton's method and a backtracking line search algorithm. Inside each Newton step, the standard HDG static condensation is performed [30], and the linear system is inverted using a sparse direct solver. The time step selection rule proceeds as follows: If the number of Newton iterations required to converge the current time step is less than or equal to the number of Newton iterations required to converge the previous step, then the time step is increased by a factor of 2. Otherwise, it is reduced by the same factor. Finally, when the time step grows beyond a given value (usually 20 convective times for the geometry of interest), a steady-state nonlinear solve is performed. Despite its simplicity, we found this time stepping rule to be robust for all the 2D cases presented here. Other techniques like local time stepping might be used in complex 3D cases in order to accelerate convergence.

3. SHOCK CAPTURING MODEL

The solution of the Euler equations for high speed flows usually contains discontinuities in the flow field in the form of shock waves or contact discontinuities. These discontinuities require a special treatment to avoid oscillations in the numerical approximation because of the appearance of Gibbs phenomenon. For the case of shocks, these oscillations need to be suppressed to maintain the stability of the computation.

In this paper, we propose an artificial viscosity approach inspired by the work of Nguyen and Peraire [27] and Premasuthan *et al.* [16], which can be traced back to ideas by Bhagatwala and Lele [17] in the LES context. This artificial viscosity model is based on a sensor that is triggered around shocks and is crafted so that the shock is spread to a thickness that is matched to the grid resolution.

The key to the present approach is to use a non-dimensional form of the divergence of the velocity, or dilation, $(\nabla \cdot \mathbf{v})$ as a point-wise sensor that drives the application of artificial viscosity ϵ to the system. The artificial viscosity enters the governing equations through an artificial dissipation term that is added to the conservation law in the following manner:

$$\frac{\partial \mathbf{u}}{\partial t} + \nabla \cdot \mathbf{F}(\mathbf{u}) = \nabla \cdot \epsilon \nabla \mathbf{u}_{\mathrm{AV}}(\mathbf{u}). \tag{4}$$

Here, $\mathbf{u}$ are the conserved variables in the Euler equations, $\mathbf{F}$ denotes the inviscid fluxes, and the terms on the right hand side represent the viscosity, which is based on a Laplacian-like term applied

DOI: 10.1002/fld

to $\mathbf{u}_{\mathrm{AV}}(\boldsymbol{u}) = \{\rho, \rho\mathbf{v}, \rho H\}$. The difference between $\mathbf{u}$ and $\mathbf{u}_{\mathrm{AV}}$ on the last term is designed to ensure conservation of enthalpy across shocks [26] for steady-state simulations while still being a dissipative term in the transient case [34].

3.1. Non-dimensional shock sensor

The amount of artificial viscosity added to the equations is given by ϵ, which is a function of the dilation, the state of the fluid, and the mesh resolution. In order to make the dependency of ϵ on $\nabla \cdot \mathbf{v}$ non-dimensional, we need to define a characteristic scale for length and velocity.

The choice of these scales is driven by the need to capture the shock within an element. In low-order schemes, the length scale is of the order of the element size, denoted by h. For high-order finite element methods, as described by Persson and Peraire [23], such length scale is of the order of h/p, where p denotes the polynomial order used to approximate the solution. In general, h is not constant in the whole domain, but rather, it varies from element to element. Therefore, it needs to be defined at every point in the domain as a scalar field $h = h(\mathbf{x})$.

The simplest approach is to assign an element size h_e to each element in the mesh and assume that h is constant within each element. In this way, the resulting element size field $h(\mathbf{x})$ becomes piecewise constant. This results in a discontinuous artificial viscosity field, which degrades the accuracy of the solution and might lead to instabilities as shown by Barter and Darmofal [26]. For this reason, we propose to use a piecewise linear reconstruction similar to the one they proposed [26] by averaging the element size, h_e, for all the elements surrounding a vertex.

This process requires the definition of h_e, which is crucial when anisotropic adaptation is performed and elements align and stretch along the discontinuities. In that instance, the natural length scale for the shock is the element dimension in the direction of the gradient of the solution. A good surrogate for this is the smallest altitude of the element, defined as the minimum distance from any vertex of the polygon (triangle or quadrilateral in 2D) to any opposite side, as depicted in Figure 1. In the case of high-order meshes, the element size field is computed before the mesh is curved.

Using this reconstruction, we define the reference length scale to be equal to $k_h h(\mathbf{x})/p$, where k_h is a correction term of $\mathcal{O}(1)$, which will be determined empirically from 1D computations.

In addition to this, the non-dimensionalization of $\nabla \cdot \mathbf{v}$ requires a velocity scale. Other authors have proposed to use the speed of sound $c = \sqrt{\gamma RT}$, which is subject to change across the shock. This variation is especially severe for strong shocks. To minimize this effect, we propose to use the critical speed of sound c^*:

$$c^* = \sqrt{\gamma R T^*} = \sqrt{\gamma R \left(\frac{2}{\gamma + 1}\right) T_0} \tag{5}$$

that only depends on the total temperature, T_0, which is constant across a stationary shock.

Using these two scales, we can define the non-dimensional shock sensor $\tilde{s}^*(\boldsymbol{u})$ as follows:

$$\tilde{s}^*(\boldsymbol{u}) = -\frac{(k_h h/p)\nabla \cdot \mathbf{v}}{c^*}. \tag{6}$$

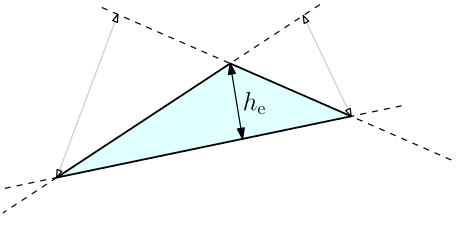

Figure 1. Sketch of the procedure to extract h_e for triangles, based on identifying the smallest altitude in the polygon. All the altitudes are colored in light gray except for the smallest one, colored in solid black. The use of the smallest altitude ensures that the correct length scale is used under the assumption that anisotropic meshes align with shocks.

DOI: 10.1002/fld

This sensor becomes positive in the presence of shock waves. Moreover, a simple analysis based on 1D shock theory shows that this sensor asymptotes to a constant as the shock strength increases. In particular, if we assume that the shock thickness is of the order of $k_h h/p$, then

$$\nabla \cdot \mathbf{v} \approx \frac{v_{2n} - v_{1n}}{k_h h/p}, \tag{7}$$

where v_{1n} and v_{2n} are the normal velocities before and after the shock, respectively. Hence,

$$\tilde{s}^* = -\frac{(k_h h/p)\nabla \cdot \mathbf{v}}{c^*} \approx -\frac{v_{2n} - v_{1n}}{c^*} = \left(1 - \frac{v_{2n}}{v_{1n}}\right) M_{1n} \frac{c_1}{c^*}, \tag{8}$$

which can be further expanded into

$$\tilde{s}^* \approx \frac{2M_{1n}^2 - 2}{(\gamma+1)M_{1n}} \sqrt{\frac{\gamma+1}{2 + (\gamma-1)M_{1n}^2}} \tag{9}$$

and asymptotes to $\tilde{s}^* = 2/\sqrt{\gamma^2 - 1}$ for very high incident Mach number, M_{1n}. This result is confirmed in Section 4 using 1D simulations. For comparison purposes, we note that the same simplified analysis applied to the case where the velocity scale is given by the speed of sound predicts a discrepancy in the sensor before, $\tilde{s}_1$, and after, $\tilde{s}_2$, the shock given by the following:

$$\frac{\tilde{s}_1}{\tilde{s}_2} \approx \frac{c_2}{c_1} = \sqrt{\frac{T_2}{T_1}} = \sqrt{\frac{\left(2\gamma M_{1n}^2 - (\gamma-1)\right)\left((\gamma-1)M_{1n}^2 + 2\right)}{(\gamma+1)^2 M_{1n}^2}} \tag{10}$$

that grows linearly with the Mach number ahead of the discontinuity. The behavior of $\tilde{s}_1$, $\tilde{s}_2$, and $\tilde{s}^*$ as a function of M_{1n} is shown in Figure 2.

3.2. *Viscosity scale*

Dimensional consistency dictates that the artificial viscosity has units of velocity times length. For the latter, we use the already defined characteristic length scale $h(\mathbf{x})$; hence, only the velocity scale requires attention. A natural choice would be the fastest wave across the shock, given by $\lambda_{\max} = |\mathbf{v} \cdot \mathbf{n}| + c$; however, this requires the extraction of the vector normal to the shock front $\mathbf{n}$. A simpler choice is to take $\sqrt{\mathbf{v} \cdot \mathbf{v} + c^2}$ instead of $\lambda_{\max}$. With this, the viscosity scale is given by the following:

$$\epsilon_{AV} = \left(k_h \frac{h}{p}\right) \sqrt{\mathbf{v} \cdot \mathbf{v} + c^2}. \tag{11}$$

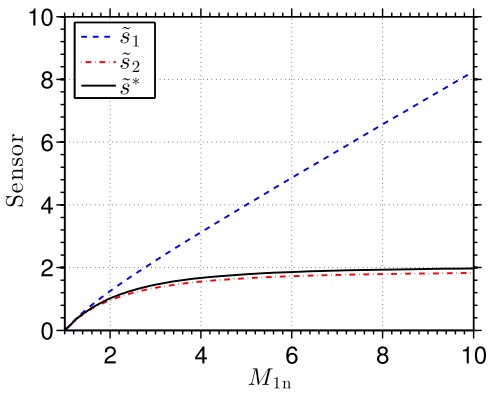

Figure 2. Estimates for $\tilde{s}_1$, $\tilde{s}_2$, and $\tilde{s}^*$ as a function of the incident Mach number M_{1n}, based on the assumption that the shock is one-dimensional and captured within the available resolution. Note the divergence between $\tilde{s}_1$ and $\tilde{s}_2$ with the shock strength and that $\tilde{s}^*$ asymptotes to a constant.

DOI: 10.1002/fld

Provided that the shock sensor is of order 1, this definition of viscosity scale yields a cell Peclet number $Pe_{\text{cell}} = \mathcal{O}(1)$, which reconciles well with linear theory for convection–diffusion and the limit of $Pe_{\text{cell}} < 2$ for oscillation-free solutions.

3.3. Artificial viscosity

To define the final form of the artificial viscosity, we follow the approach by Nguyen and Peraire [27] and introduce a soft-max function to switch on the viscosity when the indicator is positive (denoting compression or shock) and minimize its effect in regions where the flow expands. Namely, we propose to add viscosity according to the following formula:

$$\epsilon = \left(k_h \frac{h}{p}\right) \sqrt{\mathbf{v}\cdot\mathbf{v} + c^2}\; f\left(\tilde{s}^*\right), \tag{12}$$

where $\tilde{s}^*$ is defined by (6) and the function f is defined as follows:

$$f(x) = \frac{\log\left(1 + \exp\left(\alpha(x-\beta)\right)\right)}{\alpha}. \tag{13}$$

We note that this function is a smooth approximation to $\max(0, x-\beta)$ that is introduced to reduce the risk of convergence problems because of discontinuity on the slope. In addition, the following bounds hold:

$$0 < f(x) - \max(0, x-\beta) < (\log 2)/\alpha. \tag{14}$$

We note that α governs the shape of the function $f(x)$ around $x = \beta$, while β controls the kick-in value of the sensor $\tilde{s}$, which is a small positive number. This introduces a small gap in which $\tilde{s}^* > 0$ and ϵ are small compared with the numerical diffusion and is required to recover optimal asymptotic convergence in the presence of smooth compressions (Section 5.4). After some numerical tests, we propose to use $\alpha = 10^4$ and $\beta = 0.01$.

4. ONE-DIMENSIONAL STUDIES

Here, we perform some simple 1D computations in order to determine the value of k_h and assess the performance of the artificial viscosity model proposed in a simple case over a range of shock strengths and approximation orders.

For all the results presented here, we consider a stationary shock wave in a tube modeled using the 1D Euler equations. Given a Mach number upstream of the shock, the corresponding state behind the shock can be computed using standard normal shock theory. These two states provide the supersonic inflow and subsonic outflow boundary conditions. To obtain a steady-state solution to this problem, we initialize the flow to a smooth profile between the two states and march it in time using a backward Euler discretization. In principle, this problem is ill posed because of the ambiguity in location of the shock wave. We constrain the shock position by making the value of the density at the midpoint of the domain be equal to the average value across the shock [23].

4.1. Choice of length-scale correction factor

To determine the effect of k_h on the shock profile, we simulate a 1D steady shock in the domain $x \in [0, 1]$, using 40 elements of order $p = 4$. For this study, only two parameters are varied: $M_{1\text{n}} = \{2, 5\}$ and $k_h = \{0.5, 1, 1.5, 2, 3\}$. The results for all these cases are summarized in Figure 3. The x-axis in Figure 3 is transformed to non-dimensional units centered at the shock, using the element size h_{e} as reference length. In both cases, the shock profiles agree with the expected behavior in that the shock gets wider as k_h grows. One can readily identify $k_h < 1$ as an unsuitable choice due to the strong oscillations in both pressure and density, as well as the shock sensor itself. Similarly, values in the region $k_h > 2$ are not interesting as they do not show improvement over the $k_h = 2$ solution. We conclude that $k_h \in [1, 2]$ is a reasonable choice. We propose to take $k_h = 1.5$ and use this value in the following simulations.

DOI: 10.1002/fld

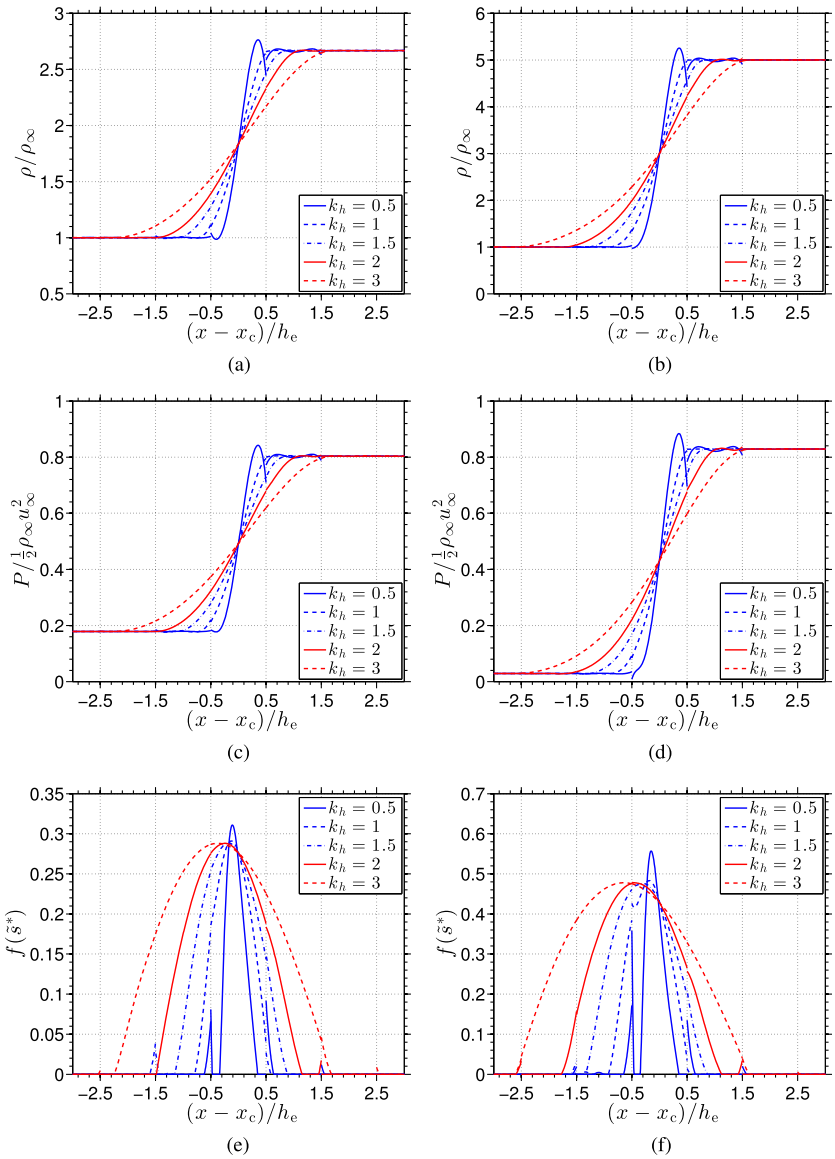

Figure 3. Density, pressure, and sensor profile around a 1D shock with $M_{1n} = \{2, 5\}$, computed using different values of $k_h = \{0.5, 1, 1.5, 2, 3\}$ and polynomial degree $p = 4$. For $k_h \in [1, 2]$, the shock profiles virtually free of oscillations. (a) Density, $M_{1n} = 2$, (b) density, $M_{1n} = 5$, (c) pressure, $M_{1n} = 2$, (d) pressure, $M_{1n} = 5$, (e) shock sensor, $M_{1n} = 2$, and (f) shock sensor, $M_{1n} = 5$.

4.2. Mach number study

Here, we present 1D simulations for a wide variation of the upstream Mach number $M_{1n} = \{2, 5, 10, 20, 30\}$. As in the previous results, the domain is composed of 40 uniform elements of order $p = 4$. The density, pressure, and shock indicator profiles are plotted in Figure 4. We observe that the shock width is independent of the strength that is because of the definition of the length and velocity scales for the sensor as well as the viscosity. These results indicate that the formulation is robust for high Mach number flows. That is, the transition between states in the shock region is insensitive to the properties of the shock and happens within the available resolution. Furthermore, the results demonstrate that the analysis outlined in Section 3.1 holds, that is, the sensor saturates as the shock gets stronger, and hence validates the choice of c^* as the velocity scale for the shock sensor.

DOI: 10.1002/fld

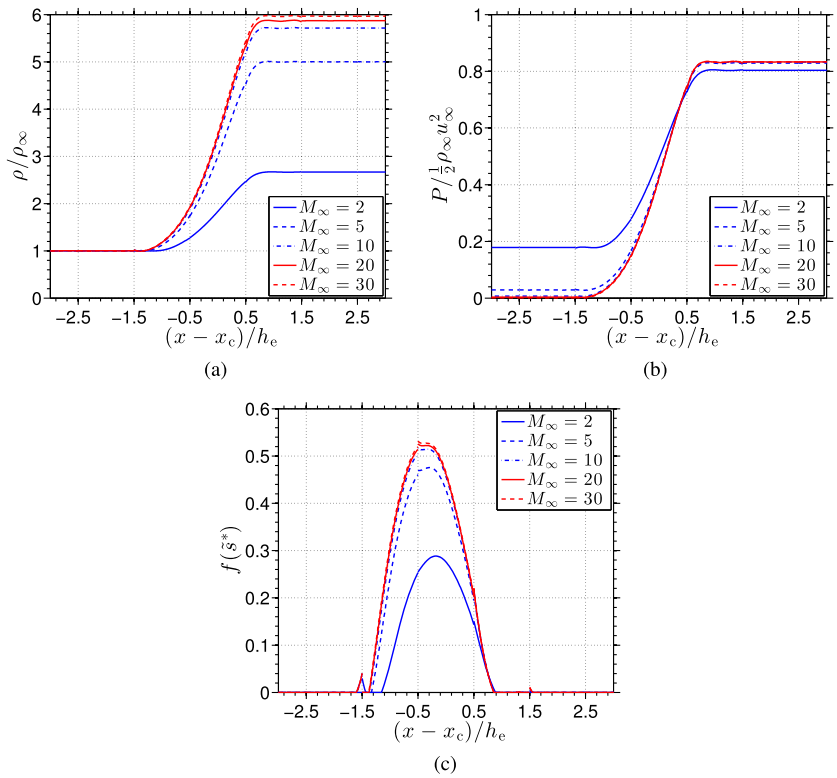

Figure 4. Density, pressure, and sensor profile computed using different values of the upstream Mach number $M_{1n} = \{2, 5, 10, 20, 30\}$ for $k_h = 1.5$ and $p = 4$. The shock sensor asymptotes to a constant profile as predicted by the analysis in Section 3.1. (a) Density, (b) pressure, and (c) shock sensor.

4.3. Polynomial order study

Here, we consider the case of $M_{1n} = 5$ and modify the polynomial order to verify the effect of the h/p scaling. For these runs, we consider the polynomial orders: $p = \{2, 3, 4, 5, 6, 7, 8\}$. The results of this experiment are compiled in Figure 5. We can readily identify a narrowing of the solution profile as the polynomial order is increased, without significant oscillations in the solution. Furthermore, the shock thickness scales like $1/p$ as expected from the definition of the viscosity scale (Equation 11). In that respect, the value $k_h \in [1, 2]$ derived using solutions with $p = 4$ seems to be independent of the approximation order.

4.4. Shock center study

All the cases shown previously require a numerical fix to avoid ill-posedness by ensuring that the value of the density at one prescribed point in the domain is equal to the average value across the shock. This removes the degree of freedom associated to the location of the shock and makes the steady-state problem solvable. For simplicity, this point has been made to coincide with the center of an element, which is not generally true in more complex cases. Next, we want to modify the constraint to allow for the shock center to be located anywhere along an element. The results obtained for the case of $M_{1n} = 5$ using polynomials of order $p = 4$ are shown in Figure 6. The results indicate that the profile is not affected by the alignment of the mesh and the shock, at least in one dimension. This property is crucial in moving to more complex cases where the exact position of the shock is unknown *a priori*. Furthermore, this behavior ensures that the scheme does not introduce spurious oscillations as the solution crosses the element boundaries in unsteady cases.

DOI: 10.1002/fld

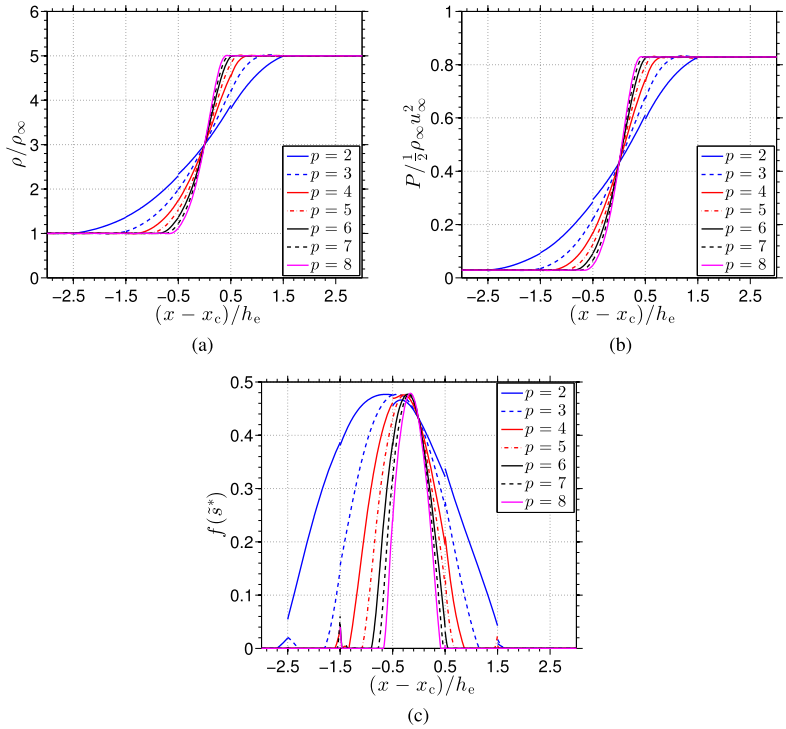

Figure 5. Density, pressure, and sensor profile computed using different polynomial orders $p = \{2, 3, 4, 5, 6, 7, 8\}$ for $M_{1n} = 5$ and $k_h = 1.5$. The results are clean of oscillations in the primal variables for all the orders of approximation. (a) Density, (b) pressure, and (c) shock sensor.

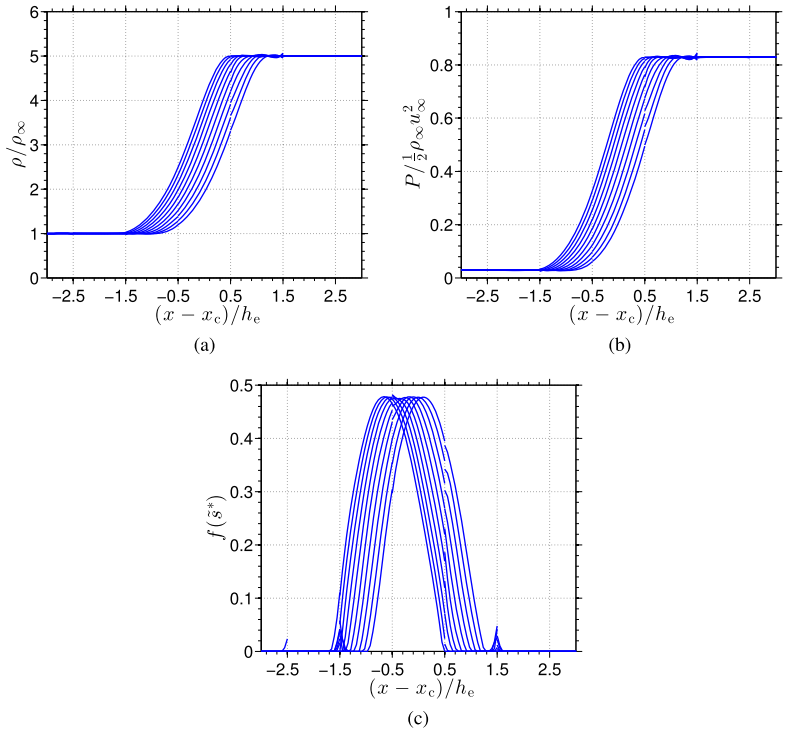

Figure 6. Density, pressure, and sensor profile computed using different locations for the shock midpoint along an element for $M_{1n} = 5$, $p = 4$, and $k_h = 1.5$. The results indicate that the shock profile is not affected by the particular location of the shock midpoint within one element. (a) Density profile, (b) pressure profile, and (c) shock sensor.

DOI: 10.1002/fld

5. TWO-DIMENSIONAL STUDIES

In this section, we focus our attention on 2D test cases. In particular, we aim to demonstrate the performance of our artificial viscosity model on highly anisotropic meshes. The use of adaptive anisotropic mesh methods introduces two challenges from a shock-capturing perspective, the ability to compute solutions on coarse meshes typical of early phases of the adaptivity and the ability to maintain sub-cell shock resolution for highly adapted anisotropic meshes towards the end of the process. To demonstrate the capability of our method in dealing with adaptive anisotropic meshes, we perform adaptive mesh refinement on the solution with the help of the BAMG code by Hecht [35].

5.1. Supersonic flow in a duct with a ramp

The first case considered is the simulation of a supersonic inviscid flow at $M_{\text{in}} = 1.5$ inside a duct with a compression ramp at 5° of inclination. In this case, there is a limited interaction between the shock wave and the expansion fan.

The simulation starts from an isotropic mesh with minor refinement depicted in Figure 7(a). At each step, a steady-state solution is computed using polynomials of order $p = 4$; then BAMG is used to generate a new mesh. For all the results generated in this paper, BAMG was run with a minimum edge size of 0.001, a maximum edge size of 0.5, and two iterations of Jacobi smoothing. The final mesh, obtained after nine iterations, is adapted to capture the shock bouncing off the walls as well as the expansion wave emanating at the end of the ramp (Figure 7(c)). In this case, the maximum aspect ratio amongst the elements of the mesh is 250. The shock sensor field is depicted in Figure 8. We note that the artificial viscosity model is robust enough to converge the problem during the early iterations of adaptivity when the mesh is coarse and the solution is under-resolved. This same model is capable of producing more accurate solutions as anisotropy appears in the mesh

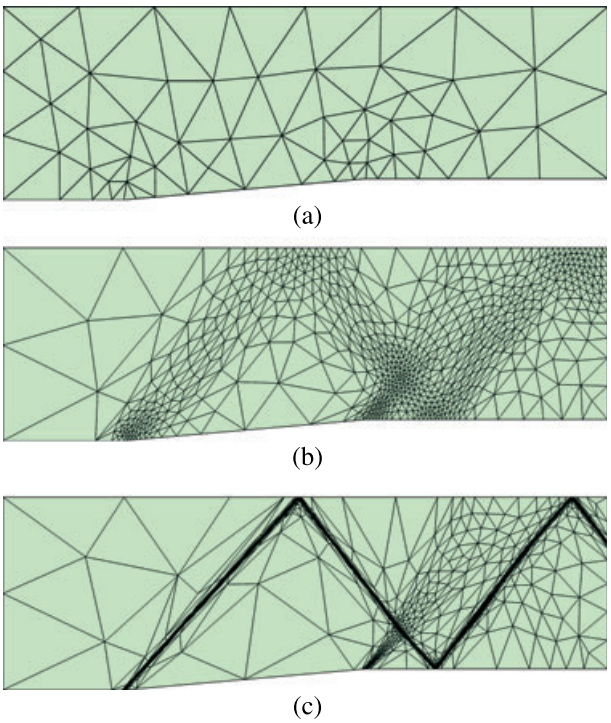

Figure 7. Evolution of the anisotropic mesh adaptivity iteration, for the case of the supersonic compression ramp. (a) Initial mesh, 102 elements, (b) adapted mesh after two iterations, 1324 elements, and (c) adapted mesh after nine iterations, 2439 elements.

DOI: 10.1002/fld

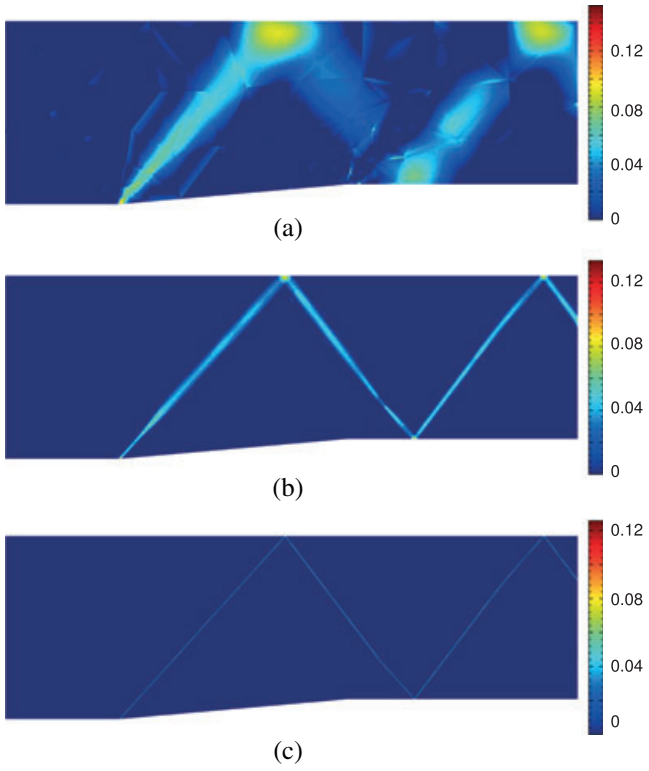

Figure 8. Evolution of the shock sensor with the anisotropic mesh adaptivity iteration. The proper scaling of the shock removes the dependency of the sensor magnitude on the grid. (a) Shock sensor on initial mesh, (b) shock sensor after two iterations of anisotropic refinement, and (c) shock sensor after nine iterations of anisotropic refinement.

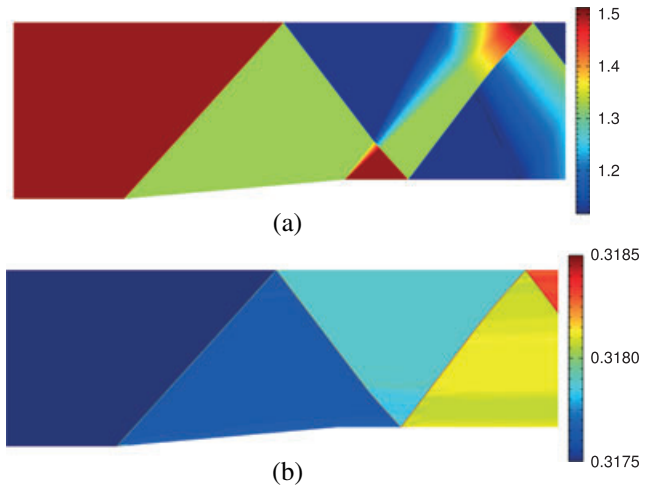

Figure 9. Mach number field (a) and entropy field (b) around the supersonic ramp computed using polynomials of order $p = 4$ on the final mesh. The combination of anisotropic adaptivity and artificial viscosity yields sharp shocks with a limited number of degrees of freedom.

and the flow is better resolved. We observe that the sensor is active over thinner regions as the mesh is refined but its magnitude does not change, thus showing mesh independence. A sample of the results obtained on the finest mesh is compiled in Figure 9. Notice the clean resolution of shocks and expansion fans that is visible on both the Mach number and the entropy in the domain.

DOI: 10.1002/fld

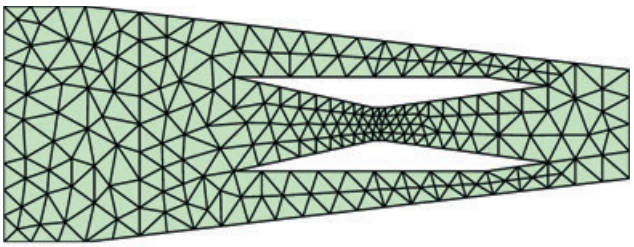

(a)

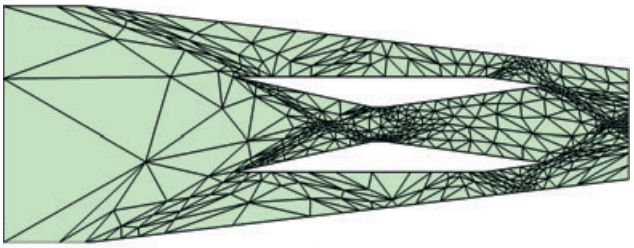

(b)

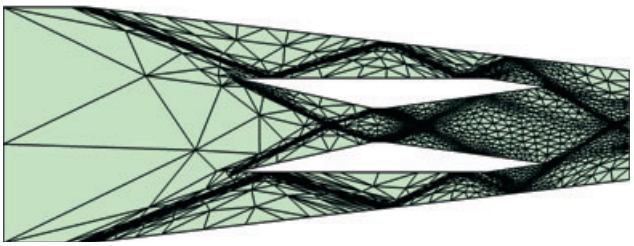

(c)

Figure 10. Evolution of the mesh for the Scramjet flow as a function of the number of iterations of anisotropic adaptivity. (a) Initial mesh, 461 elements, (b) adapted mesh after four cycles, 898 elements, and (c) adapted mesh after 12 cycles, 6509 elements.

5.2. Supersonic flow inside a Scramjet geometry

In this case, we consider the solution of the flow at $M_\infty = 3.6$ inside a Scramjet modeled using the 2D surrogate geometry proposed by Kumar [36]. This geometry is composed of two inner bodies inside a duct, all of which are straight sided. The problem is symmetric across the horizontal axis although this is not enforced in the meshes we use. This configuration produces a variety of flow features (shock waves, expansion waves, and contact discontinuities) that interact with each other to form a complex pattern.

As in the previous case, we start from an isotropic coarse mesh with moderate refinement at the corners. At every step, we compute a steady-state solution that is used as a basis for BAMG to generate a new mesh. All the intermediate solutions required in this process are computed using polynomials of order $p = 4$. A total of 12 adaptivity iterations are performed. The last iteration yields a mesh with a maximum aspect ratio of 125. A sample of the mesh evolution is summarized in Figure 10. The flow field computed on the final mesh is depicted in Figure 11.

We verify that the flow field remains symmetric despite the fact that symmetry is not enforced explicitly and the mesh is not symmetric. This indicates a certain level of robustness of the method to mesh misalignments. We also observe how the artificial viscosity is only targeting the shocks and not other flow features like contacts discontinuities or expansion waves. One place where this is visible is the region around the wall in the nozzle throat. A detail of the flow there is shown in Figure 12. We also note how the shock sensor is only active in the shock wave and is turned off at the intersection with the expansion fan. Something similar happens at the trailing edge of the inner body, where the shock sensor is not activated by the presence of a slip line as shown in Figure 13.

DOI: 10.1002/fld

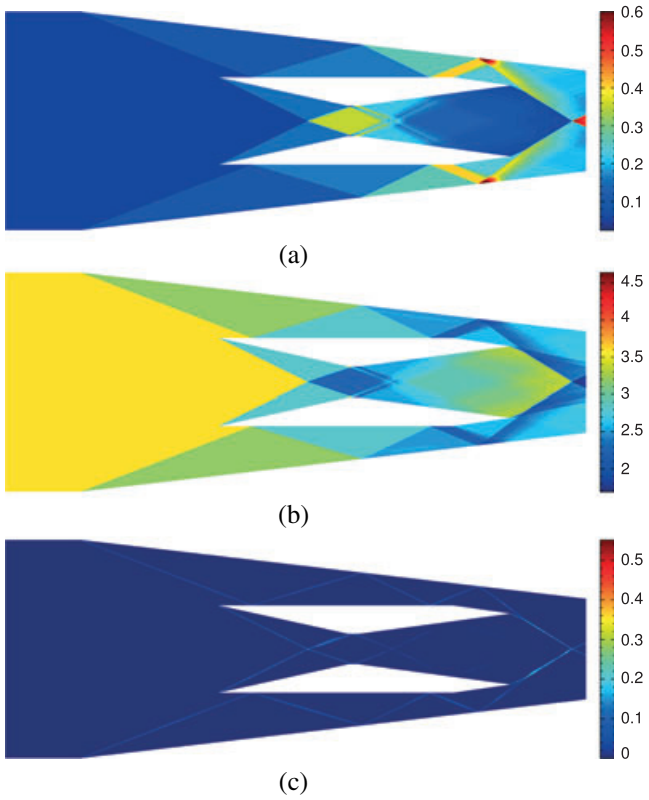

Figure 11. Flow inside the Scramjet geometry computed on the finest mesh (12 adaptivity cycles) using polynomials of order $p = 4$. The combination of adaptivity and the sub-cell shock resolution of the proposed scheme can separate very fine details in the flow. (a) Pressure, (b) Mach number, and (c) shock sensor.

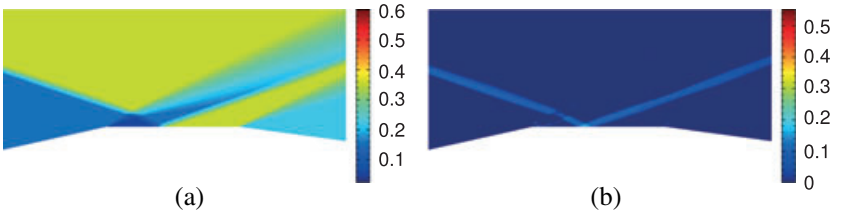

Figure 12. Detail of the flow around the nozzle throat. The artificial viscosity is only applied in the shock region and is turned off at the intersection with the expansion fan. (a) Pressure and (b) shock sensor.

5.3. Supersonic flow over a forward facing step

Here, we use the proposed shock-capturing scheme to simulate the well-known case proposed by Woodward and Collela [37] consisting of a supersonic flow at $M_\infty = 3$ inside a straight channel with a forward facing step. This problem was designed to test a variety of finite volume schemes and since then has been used as a benchmark for compressible flow solvers. Our interest here lies in the steady-state version of the problem that contains rich physics that will further validate the artificial viscosity model.

In order to avoid the strong singularity caused by the expansion at the corner of the step, we modify the geometry by replacing the corner with quarter circle of radius 1% of the total height of the step. This modification is similar in spirit (although not in form) to the modified discretization that Woodward and Collela used in the corner [37], or the viscosity that some schemes based on a resolution indicator add there [25, 38].

DOI: 10.1002/fld

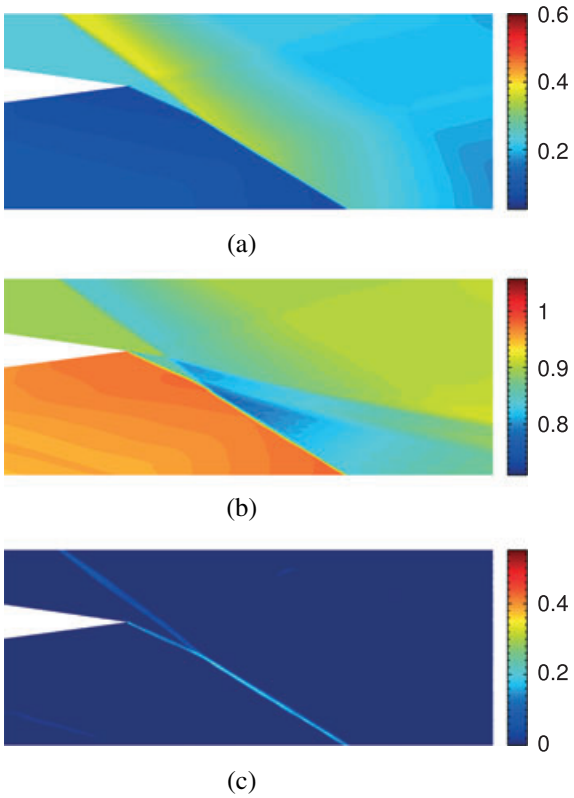

Figure 13. Detail of the flow around the trailing edge of the inner body. The presence of a contact discontinuity, which is a purely vortical feature, does not trigger the shock sensor. (a) Pressure, (b) magnitude of the velocity, and (c) shock sensor.

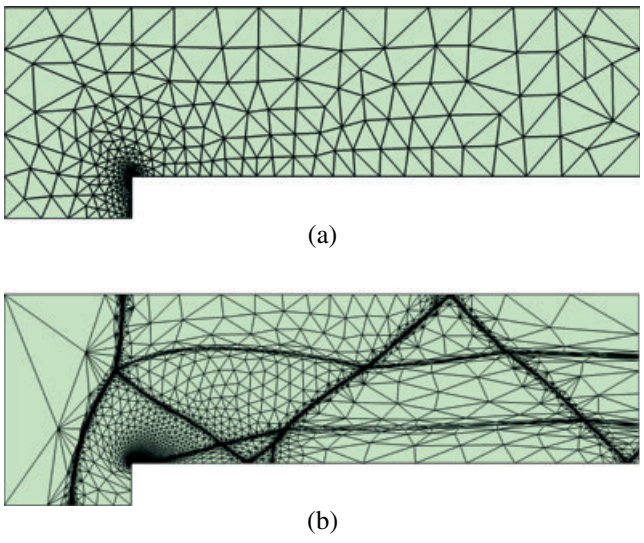

Figure 14. Evolution of the mesh for the forward facing step flow as a function of the number of anisotropic adaptivity iterations. (a) Initial mesh, 677 elements and (b) adapted mesh after 12 cycles, 9584 elements.

We compute the solution to this flow on a series of meshes generated with BAMG using polynomials of order $p = 4$. In total, 12 cycles of adaptation are enough to produce the final mesh shown in Figure 14. In this case, the maximum aspect ratio amongst the elements of the mesh is 70. The initial mesh is also plotted in the same figure. The solution field for the pressure, Mach number, and

DOI: 10.1002/fld

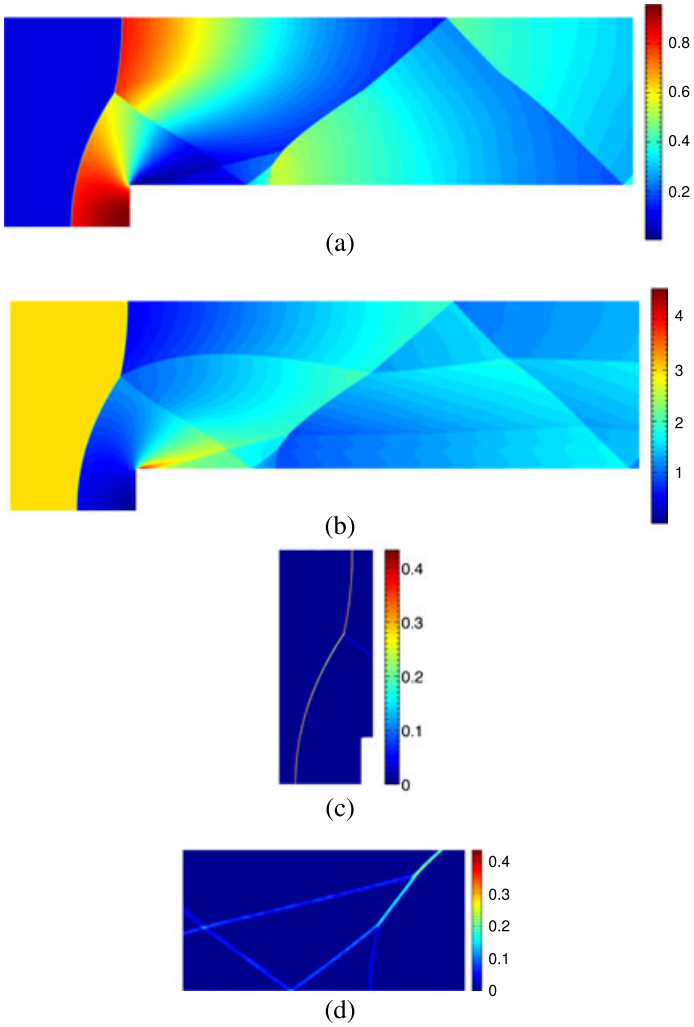

Figure 15. Flow solution for the forward facing step problem computed using polynomials of order $p = 4$ on the finest mesh (after 12 refinement iterations). Sharp flow features (e.g., shock waves, slip lines, etc.) are properly captured by the mesh adaptivity and the artificial viscosity model, while smooth regions benefit from the high-order approximation. (a) Pressure, (b) Mach number, (c) detail of the shock sensor around the bow shock, and (d) detail of the shock sensor at the wall.

shock sensor is plotted in Figure 15. We note how the different flow features are properly described within the available mesh resolution. In particular, we can readily observe the presence of a triple point in the bow shock that leads to a contact discontinuity as well as a reflected shock. The latter interacts with the expansion around the corner to generate a complex flow pattern that includes a weak normal shock, and two additional triple points. Two out of the three contact discontinuities are picked up by the anisotropic refinement along the adaptation cycle and can be seen in the final meshes (Figure 14(b)). However, the weakest one, associated to the merging of the reflected shock and the weak normal shock close to the wall, is not. This problem depends on the adaptivity strategy and could be fixed by carefully tuning the parameters that govern BAMG but was considered outside the scope of this paper.

5.4. Effect of the artificial viscosity on the convergence of the subsonic flow over a smooth bump

In this test case, we want to assess the effect of the proposed artificial viscosity model in the smooth regions of the flow (which generally represent the majority of the domain). For this, we consider the C1.1 case of the 1st International Workshop on High-Order CFD Methods [39]. This test case

DOI: 10.1002/fld

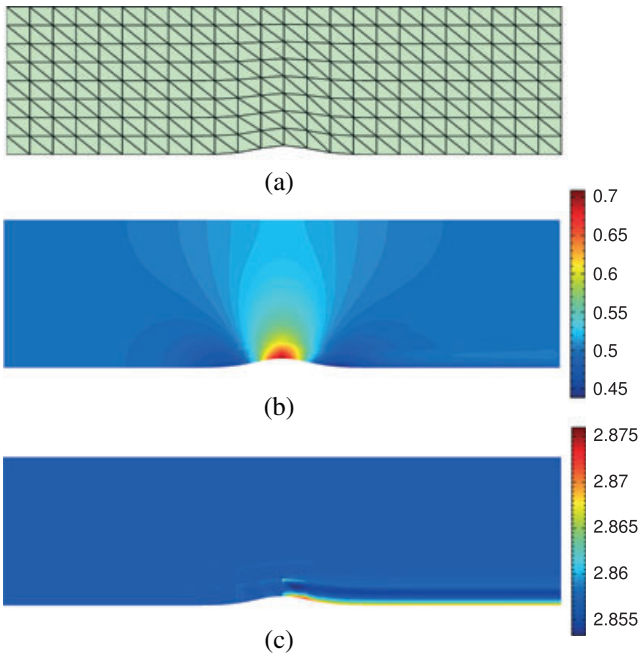

Figure 16. Sample mesh and solution for the subsonic flow over a smooth bump. The numerical solution shows oscillations in the entropy around the bump that can be quantified to assess the effect of the artificial viscosity on a smooth flow. (a) Mesh, (b) Mach number, and (c) entropy.

consists of a subsonic inviscid flow in a channel with a smooth bump on the lower surface. The entropy for the exact solution is constant. However, the numerical approximation tends to introduce small oscillations in the entropy close to the lower wall where there is flow curvature. Once generated, these oscillations persist downstream and induce a non-negligible error in the entropy norm, which is measured as follows:

$$||s - s_\infty||_2 = \sqrt{\frac{\int_\Omega \left(1 - \frac{P/\rho^\gamma}{P_\infty/\rho_\infty^\gamma}\right)^2 dV}{\int_\Omega dV}}. \tag{15}$$

We are interested in quantifying how the artificial viscosity affects this error. For this, the solution on a sequence of structured iso-parametric meshes using polynomials of order $p = \{1, 2, 3, 4\}$ is performed. For comparison purposes, we run the solver with and without the artificial viscosity terms.

The mesh generation at each stage of refinement starts from a structured mesh of quads for the bounding rectangle that is broken along the SE-NW diagonal to generate a structured mesh of triangles. The high-order nodes are placed on the rectangular mesh and then transformed analytically as described in the case specification [39]. We construct a sequence of meshes starting from a 6×2 mesh by doubling the number of elements in each direction at each stage. We repeat this process seven times for $p = 1$ and $p = 2$ and six times for the higher order cases. This is well beyond the resolution requirements for a smooth flow like this one; however, the purpose here is to grid converge the solution to assess convergence rates.

The solution on each mesh is computed using the standard HDG solver. For the cases with artificial viscosity, the standard parameters are used, namely, $\alpha = 10^4$, $\beta = 0.01$, and $k_h = 1.5$. A sample mesh and solution (Mach number and entropy) are shown in Figure 16. The solution exhibits oscillations of the entropy in the lower wall. The convergence of the entropy error for the mesh sequences is plotted in Figure 17. The horizontal scale measures the square root of the number of elements in the mesh, which is equivalent to a measure of the element size. The solid lines denote results computed using artificial viscosity, while the dashed lines denote results without it.

DOI: 10.1002/fld

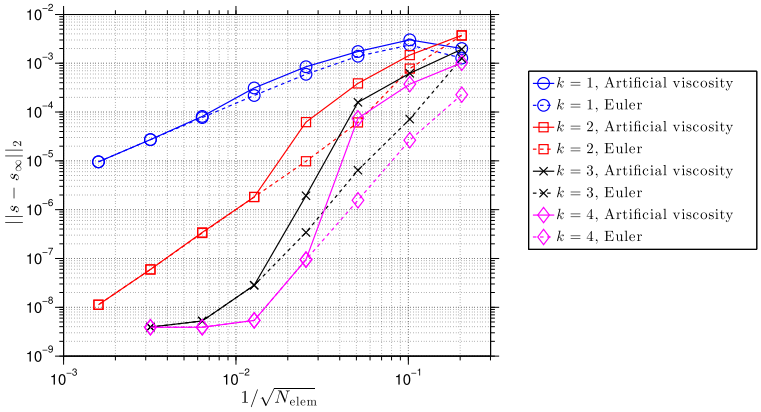

Figure 17. Comparison of the convergence in the entropy error for different polynomial orders using structured triangular meshes. The results using artificial viscosity show a pre-asymptotic behavior that seems to limit convergence between second and third orders. This effect disappears once the element size is small enough.

We can readily identify the effect of the artificial viscosity for the coarse meshes (right half of the plot) especially for the highest order polynomials. In particular, for the case of $p = 1$, the effect of adding an artificial viscosity term of order h^2 barely affects the entropy error. This is consistent with the assumption that the equivalent dissipation of the original HDG scheme is order h^2 or lower for $p = 1$, and greater than h^2 for $p \geqslant 2$.

As the mesh is refined, there is a point in which the gap introduced by the choice of α and β (Section 3.3) takes effect and the results with and without artificial viscosity coincide. This is because the flow might contain smooth regions where dilation is negative but finite. In these regions, $\tilde{s}^*$ decreases as the mesh is refined, and, at some point, is effectively turned off. This is clearly seen on the left half of Figure 17. A similar pre-asymptotic behavior was documented by Barter and Darmofal [26] for another bump geometry using a variety of shock indicators with and without their PDE smoothing.

This study indicates that the proposed scheme converges to the inviscid solution for flows without discontinuities. For practical cases, however, there is a pre-asymptotic range that is governed by the finite value of $\nabla \cdot \mathbf{v}$ in certain parts of the flow.

6. CONCLUSIONS AND FUTURE WORK

We have presented a shock-capturing strategy for high-order methods based on artificial viscosity. We derived the model, paying special attention to the scaling of the divergence of the velocity that acts as a sensor for the presence of shocks. We have tested the model extensively in 1D and 2D in combination with anisotropic adaptive refinement. The results indicate that the model can capture shocks in a variety of flows and does not interact with other flow features such as expansion fans or contact discontinuities, which do not require stabilization.

Compared with other existing models, the proposed model has several attractive features. The model does not require the solution of an extra equation to solve for the artificial viscosity field [26] because it produces artificial viscosity fields that are $\mathcal{C}^0$ in the continuous limit (and almost so in practical applications). This satisfies the conditions laid out by Barter and Darmofal [26] without the need for constructions that either limit the CFL number when treated explicitly, or widen the stencil when treated implicitly [24, 25]. Indeed, the model is fully analytical and can be implemented in any solver where gradients of the solution are available or can be computed. This is crucial to deliver steady-state solutions at a reasonable cost. Of course, the model also has its weaknesses. For example, the need for gradients of the solution that can be of low quality for certain discretization schemes and incur an extra cost in the computation. Neither of these is the case for the HDG scheme, which is ideally suited for the proposed model.

DOI: 10.1002/fld

The model could be improved in a variety of ways. First, other definitions of $\mathbf{u}_{AV}$ could yield sharper shocks or a better behavior at high Mach number [23]. Also, the whole definition of the shock length scale could be revisited to try other reconstructions for the h field as well as more sophisticated approaches based on a tensor field rather than a scalar one. The later would make sure that the dilation is measured in the metric induced by the element so that misalignments with the shock could be better identified and stabilized. This might be a requirement for the applicability to viscous flows where typical meshes present other sources of anisotropy. All these subjects will be the focus of future research.

ACKNOWLEDGEMENTS

The authors would like to acknowledge Dr. Marshall Galbraith, Dr. Steve Allmaras, and Professor David Darmofal for their feedback and discussion of various aspects of the model. The authors also acknowledge the Air Force Office of Scientific Research for supporting this work under the AFOSR grant FA9550-12-1-0357. D. Moro would like to acknowledge the financial support of the Fundacion La Caixa and the Zakhartchenko Fellowship.

REFERENCES

1. Harten A, Engquist B, Osher S, Chakravarthy SR. Uniformly high order accurate essentially non-oscillatory schemes, III. *Journal of Computational Physics* 1987; **71**:231–303.
2. Liu XD, Osher S, Chan T. Weighted essentially non-oscillatory schemes. *Journal of Computational Physics* 1994; **115**:200–212.
3. Hu C, Shu CW. Weighted essentially non-oscillatory schemes on triangular meshes. *Journal of Computational Physics* 1999; **150**(1):97–127.
4. Cockburn B, Shu CW. The Runge–Kutta local projection P1-discontinuous-Galerkin finite element method for scalar conservation laws. *Technical Report 388*, IMA, University of Minnesota: Minneapolis, MN, 1988.
5. Cockburn B, Shu CW. The Runge–Kutta discontinuous Galerkin method for conservation laws V: multidimensional systems. *Journal of Computational Physics* 1998; **141**(2):199–224.
6. Krivodonova L, Xin J, Remacle JF, Chevaugeon N, Flaherty JE. Shock detection and limiting with discontinuous Galerkin methods for hyperbolic conservation laws. *Applied Numerical Mathematics* 2004; **48**:323–338.
7. Krivodonova L. Limiters for high-order discontinuous Galerkin methods. *Journal of Computational Physics* 2007; **226**(1):879–896.
8. Von Neumann J, Richtmyer RD. A method for the numerical calculation of hydrodynamic shocks. *Journal of Applied Physics* 1950; **21**:232–237.
9. Jameson A. Analysis and design of numerical schemes for gas dynamics 1: artificial diffusion, upwind biasing, limiters and their effect on accuracy and multigrid convergence. *International Journal of Computational Fluid Dynamics* 1994; **4**:171–218.
10. Jameson A. Analysis and design of numerical schemes for gas dynamics, 2: artificial diffusion and discrete shock structure. *International Journal of Computational Fluid Dynamics* 1995; **5**:1–38.
11. Hughes TJR, Mallet M, Mizukami A. A new finite element formulation for computational fluid dynamics: II. Beyond SUPG*. *Computer Methods in Applied Mechanics and Engineering* 1986; **54**:341–355.
12. Tadmor E. Shock capturing by the spectral viscosity method. *Computer Methods in Applied Mechanics and Engineering* 1990; **80**:197–208.
13. Cook AW, Cabot WH. A high-wavenumber viscosity for high-resolution numerical methods. *Journal of Computational Physics* 2004; **195**(2):594–601.
14. Cook AW, Cabot WH. Hyperviscosity for shock-turbulence interactions. *Journal of Computational Physics* 2005; **203**(2):379–385.
15. Premasuthan S, Liang C, Jameson A. Computation of flows with shocks using the spectral difference method with artificial viscosity, II: modified formulation with local mesh refinement. *Computers & Fluids* 2014; **98**:122–133.
16. Premasuthan S, Liang C, Jameson A. Computation of flows with shocks using the spectral difference method with artificial viscosity: Part I. *Computers & Fluids* 2013; **98**:111–121.
17. Bhagatwala A, Lele SK. A modified artificial viscosity approach for compressible turbulence simulations. *Journal of Computational Physics* 2009; **228**(14):4965–4969.
18. Bassi F, Crivellini A. High-order discontinuous Galerkin discretization of transonic turbulent flows. *47th AIAA Aerospace Sciences Meeting*, Orlando, FL, 2009; AIAA 2009–180.
19. Baumann CE, Oden JT. An adaptive-order discontinuous Galerkin method for the solution of the Euler equations of gas dynamics. *International Journal for Numerical Methods in Engineering* 2000; **73**:61–73.
20. Hartmann R, Houston P. Adaptive discontinuous Galerkin finite element methods for the compressible Euler equations. *Journal of Computational Physics* 2002; **183**(2):508–532.
21. Hartmann R. Adaptive discontinuous Galerkin methods with shock-capturing for the compressible Navier–Stokes equations. *International Journal for Numerical Methods in Fluids* 2006; **51**:1131–1156.

DOI: 10.1002/fld

22. Hartmann R. Higher-order and adaptive discontinuous Galerkin methods with shock-capturing applied to transonic turbulent delta wing flow. *International Journal for Numerical Methods in Fluids* 2013; **72**:883–894.
23. Persson PO, Peraire J. Sub-cell shock capturing for discontinuous Galerkin methods. *44th AIAA Aerospace Sciences Meeting And Exhibit*, Reno, NV, 2006; AIAA 2006–112.
24. Klöckner A, Warburton T, Hesthaven JS. Viscous shock capturing in a time-explicit discontinuous Galerkin method. *Mathematical Modelling of Natural Phenomena* 2011; **10**(10):1–27.
25. Persson PO. Shock capturing for high-order discontinuous Galerkin simulation of transient flow problems. *21st AIAA Computational Fluid Dynamics Conference*, San Diego, CA, 2013; AIAA 2013–3061.
26. Barter GE, Darmofal DL. Shock capturing with PDE-based artificial viscosity for DGFEM: Part I. Formulation. *Journal of Computational Physics* 2010; **229**(5):1810–1827.
27. Nguyen NC, Peraire J. An adaptive shock-capturing HDG method for compressible flows. *20th AIAA Computational Fluid Dynamics Conference*, Honolulu, HI, 2011; AIAA 2011–3060.
28. Cockburn B, Gopalakrishnan J, Lazarov R. Unified hybridization of discontinuous Galerkin, mixed and continuous Galerkin methods for second order elliptic problems. *SIAM Journal on Numerical Analysis* 2009; **47**(2):1319–1365.
29. Moro D, Nguyen NC, Peraire J. Navier–Stokes solution using hybridizable discontinuous Galerkin methods. *20th AIAA Computational Fluid Dynamics Conference*, Honolulu, HI, 2011; AIAA 2011–3407.
30. Peraire J, Nguyen NC, Cockburn B. A hybridizable discontinuous Galerkin method for the compressible Euler and Navier–Stokes equations. *48th AIAA Aerospace Sciences Meeting*, Orlando, FL, 2010; AIAA 2010–363.
31. Nguyen NC, Peraire J. Hybridizable discontinuous Galerkin methods for partial differential equations in continuum mechanics. *Journal of Computational Physics* 2012; **231**(18):5955–5988.
32. Huerta A, Angeloski A, Roca X, Peraire J. Efficiency of high-order elements for continuous and discontinuous Galerkin methods. *International Journal for Numerical Methods in Engineering* 2013; **96**:529–560.
33. Nguyen NC, Peraire J, Cockburn B. An implicit high-order hybridizable discontinuous Galerkin method for the incompressible Navier–Stokes equations. *Journal of Computational Physics* 2011; **230**(4):1147–1170.
34. Drela M, Merchant A, Peraire J. Elimination of spurious loss in Euler equation computations. *AIAA Journal* 2000; **38**(3):411–417.
35. Hecht F. BAMG : bidimensional anisotropic mesh generator. *Technical Report*, INRIA-Rocquencourt, 2006.
36. Kumar A, Tiwari SN. Analysis of the Scramjet inlet flow field using two-dimensional Navier–Stokes equations. *Technical Report June*, NASA Langley Research Center: Hampton, VA, 1982.
37. Woodward P, Colella P. The numerical simulation of two-dimensional fluid flow with strong shocks. *Journal of Computational Physics* 1984; **54**(1):115–173.
38. Huerta A, Casoni E, Peraire J. A simple shock-capturing technique for high-order discontinuous Galerkin methods. *International Journal for Numerical Methods in Fluids* 2012; **69**(10):1614–1632.
39. Wang ZJ, Fidkowski KJ, Bassi F, Caraeni D, Cary A, Deconinck H, Hartmann R, Hillewaert K, Huynh HT, Kroll N, May G, Persson PO, Van Leer B, Visbal MR. High-Order CFD methods: current status and perspective. *International Journal for Numerical Methods in Fluids* 2012; **72**(8):811–845.

DOI: 10.1002/fld