Journal of Computational Physics 229 (2010) 1810–1827

Contents lists available at ScienceDirect

Journal of Computational Physics

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

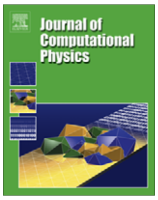

Shock capturing with PDE-based artificial viscosity for DGFEM: Part I. Formulation

Garrett E. Barter *, David L. Darmofal

Aerospace Computational Design Laboratory, Massachusetts Institute of Technology, Building 33, Room 207, Cambridge, MA 02139, USA

ARTICLE INFO

Article history:
Received 9 August 2009
Received in revised form 29 October 2009
Accepted 7 November 2009
Available online 17 November 2009

Keywords:
Discontinuous Galerkin
Compressible Navier–Stokes
Shock capturing
Artificial viscosity

ABSTRACT

Artificial viscosity can be combined with a higher-order discontinuous Galerkin finite element discretization to resolve a shock layer within a single cell. However, when a non-smooth artificial viscosity model is employed with an otherwise higher-order approximation, element-to-element variations induce oscillations in state gradients and pollute the downstream flow. To alleviate these difficulties, this work proposes a higher-order, state-based artificial viscosity with an associated governing partial differential equation (PDE). In the governing PDE, a shock indicator acts as a forcing term while grid-based diffusion is added to smooth the resulting artificial viscosity. When applied to heat transfer prediction on unstructured meshes in hypersonic flows, the PDE-based artificial viscosity is less susceptible to errors introduced by grid edges oblique to captured shocks and boundary layers, thereby enabling accurate heat transfer predictions.

1. Introduction

The focus of this work is shock capturing for higher-order discretizations with a particular emphasis on shock capturing for strong, steady shocks. A specific motivation is the simulation of hypersonic flows for the estimation of heat transfer loads. Unstructured meshes offer significant promise for computational aerothermodynamic analysis on complex geometries, but solution quality using unstructured meshes for current state-of-the art methods is far inferior to that of structured meshes [2,4,3,5,6,1]. The poor solution quality manifests itself even in symmetric, simplified test cases with poor prediction of peak heat transfer rates and asymmetric surface heat transfer distributions. The problem stems from the dependence of current shock capturing methods on the alignment of the mesh with the shocks. For simple problems, structured meshes can be designed to align the mesh with the shock such that numerical errors in shock capturing can be significantly reduced. However, for general meshes that do not align with the shock, the numerical errors at strong shocks can be significant and lead to non-physical variations that convect downstream to the boundary layer and corrupt surface heat transfer predictions.

Artificial viscosity, pioneered by von Neumann and Richtmyer [7], has been a common method of shock capturing in the context of streamwise upwind Petrov–Galerkin (SUPG) finite element methods, as proposed by Hughes et al. [8–11]. Researchers such as Hartmann and Houston [12,13] and Aliabadi et al. [14] have adopted this approach for use in discontinuous Galerkin methods as well, with good results, albeit only for $p = 1$ polynomial solutions.

The approach developed in this work is an extension of the sub-cell shock capturing method proposed by Persson and Peraire [15]. Specifically, Persson and Peraire introduced an elementwise-constant artificial viscosity that scales with the resolution length scale of a higher-order finite element method, h/p, such that the shock width is also $\mathcal{O}(h/p)$. Thus, for sufficiently high p, the shock can be captured within a single element. To locate the shocks in the flow field, Persson and Peraire

---

* Corresponding author.
E-mail addresses: gbarter@alum.mit.edu (G.E. Barter), darmofal@mit.edu (D.L. Darmofal).

0021-9991/$ - see front matter © 2009 Elsevier Inc. All rights reserved.
doi:10.1016/j.jcp.2009.11.010

developed a sensor based on the magnitude of the highest-order coefficients in an orthonormal representation of the solution.

As will be described in Section 3, an elementwise-constant artificial viscosity model has some inherent shortcomings. Specifically, element-to-element jumps in artificial viscosity lead to oscillations in state gradients that can corrupt the smoothness and accuracy of the downstream flow field. This research develops a smoother artificial viscosity by employing an artificial viscosity PDE model which is appended to the system of governing equations. Further, by combining the smooth PDE-based artificial viscosity with higher-order discretizations, the dependence of shock capturing on the grid orientation can be significantly reduced.

While the discretization used in this paper is the discontinuous Galerkin (DG) finite element method (FEM), the proposed shock-capturing approach is not strongly dependent on the particular higher-order method and will likely be equally effective for other discretizations in which sub-cell resolution is possible (SUPG/GLS, spectral volume, spectral difference, etc.). DG was first introduced by Reed and Hill [16] for the neutron transport equation. Much later, a foundation for DG methods applied to non-linear hyperbolic problems was established by Karniadakis et al. [17–25]. Independently, Allmaras and Giles [26,27] developed a second-order DG scheme for the Euler equations, building off of the work of van Leer [28–31]. Bassi and Rebay and Bey and Oden notably demonstrated the capabilities of DG for both the Euler and Navier–Stokes equations (including Reynolds Averaged Navier–Stokes) [32–36].

Section 2 reviews the DG FEM discretization of the compressible Navier–Stokes equations. Included in the review is the modification to the governing equations to append an artificial viscosity matrix for shock capturing. Section 3 motivates the use of a smooth, higher-order representation of artificial viscosity by highlighting the difficulties of a non-smooth formulation in one and two dimensions. Section 4 then presents the chief innovation of this research, a PDE for the control of artificial viscosity, and provides additional comparisons to a non-smooth formulation. Numerical results, including a hypersonic application of the new artificial viscosity model, are presented in Section 5.

2. Discretization

Let $\mathbf{u}(\mathbf{x},t) : \mathbb{R}^d \times \mathbb{R}^+ \to \mathbb{R}^m$ be the vector of m-state variables in d-dimensions for a general conservation law in the physical domain, $\Omega \subset \mathbb{R}^d$, given in the strong form by,

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

where $\mathcal{F}(\mathbf{u}) : \mathbb{R}^m \to \mathbb{R}^{m\times d}$ is the inviscid flux vector and $\mathcal{F}_v(\mathbf{u}, \nabla\mathbf{u}) : \mathbb{R}^m \times \mathbb{R}^{m\times d} \to \mathbb{R}^{m\times d}$ is the viscous flux.

2.1. Compressible Navier–Stokes equations

In the compressible Navier–Stokes equations, the conservative state vector is, $\mathbf{u} = [\rho, \rho v_i, \rho E]^T$, where ρ is the density, v_i is the velocity in the ith coordinate direction and E is the total internal energy. The flux vectors are,

$$\mathcal{F}_i(\mathbf{u}) = \begin{bmatrix} \rho v_i \\ \rho v_i v_j + \delta_{ij} p \\ \rho v_i H \end{bmatrix}; \quad \mathcal{F}_i^v(\mathbf{u}, \nabla\mathbf{u}) = \begin{bmatrix} 0 \\ \tau_{ij} \\ v_j \tau_{ij} + \kappa_T \frac{\partial T}{\partial x_i} \end{bmatrix},$$

where p is the static pressure, $H = E + p/\rho$ is the total enthalpy, δ_{ij} is the Kronecker delta, τ_{ij} is the shear stress defined below, κ_T is the thermal conductivity, $T = p/\rho R$ is the temperature and R is the gas constant. The pressure is related to the state vector by the equation of state, $p = (\gamma - 1)\rho(E - 0.5 v_i v_i)$, where γ is the ratio of specific heats. The shear stress is, $\tau_{ij} = \mu(\partial v_i/\partial x_j + \partial v_j/\partial x_i) - \delta_{ij}\lambda \partial v_k/\partial x_k$, where μ is the dynamic viscosity and $\lambda = -\frac{2}{3}\mu$ is the bulk viscosity coefficient. Here the dynamic viscosity is assumed to adhere to Sutherland's Law, and the thermal conductivity is related to the viscosity by the Prandtl number, Pr. In the remainder of the paper, the following notation will be used for the viscous fluxes,

$$\mathcal{F}_i^v(\mathbf{u}, \nabla\mathbf{u}) = \mathbf{A}_{ij}^v(\mathbf{u}) \frac{\partial \mathbf{u}}{\partial x_j},$$

where $\mathbf{A}_{ij}^v \in \mathbb{R}^{m\times m}$.

The discontinuous Galerkin finite element discretization proceeds by deriving a weak form of (1). The domain is subdivided by a triangulation, $\mathcal{T}_H$, into a set of non-overlapping elements, κ, such that $\Omega = \bigcup_{\kappa\in\mathcal{T}_H} \kappa$. Then, define $\mathcal{V}_H^p$, a vector-valued function space of discontinuous, piecewise-polynomials of degree p,

$$\mathcal{V}_H^p \equiv \{ v \in L^2(\Omega) | v|_\kappa \in P^p, \ \forall \kappa \in T_H \}.$$

The discontinuous Galerkin formulation is obtained by multiplying (1) by a test function, $\mathbf{v}_H \in (\mathcal{V}_H^p)^m$, integrating by parts, and accounting for the jumps from element-to-element by carefully defining the inviscid and viscous fluxes on element boundaries. The resulting DG formulation can then be expressed as the solution $\mathbf{u}_H(\cdot,t) \in (\mathcal{V}_H^p)^m$ to the semi-linear weighted residual (linear in the second argument), $\mathcal{R}(\mathbf{u}_H, \mathbf{v}_H) = 0, \ \forall \mathbf{v}_H \in (\mathcal{V}_H^p)^m$.

The inviscid flux functions utilized in this work are described in more detail in the remainder of the paper. The viscous flux contributions are discretized according to the second form of Bassi and Rebay [34]. The specific implementation of boundary conditions for the Navier–Stokes equations can be found in Oliver [37] and Fidkowski et al. [38]. Additional details of the entire discretization can also be found in Barter [39].

2.2. Artificial viscosity matrix

When artificial viscosity is added to the system for the purposes of resolving discontinuities, the viscous flux is augmented such that

$$\mathcal{F}_i^v(\mathbf{u}, \nabla\mathbf{u}) = \left(\mathbf{A}_{ij}^v + \mathbf{A}_{ij}^\epsilon\right)\frac{\partial \mathbf{u}}{\partial x_j},$$

A simple, diagonal artificial viscosity matrix that accounts for anisotropic meshes is given by,

$$\mathbf{A}_{ij}^\epsilon = \hat{\epsilon}(\epsilon)\mathrm{diag}\left[\mathcal{I}\frac{h_i}{\bar{h}}\right]\delta_{ij}, \tag{2}$$

$$\hat{\epsilon} \equiv \begin{cases} 0, & \epsilon \leqslant \theta_L, \\ \frac{1}{2}\theta_H\left(\sin\left[\pi\left(\frac{\epsilon-\theta_L}{\theta_H-\theta_L}-\frac{1}{2}\right)\right]+1\right), & \theta_L < \epsilon < \theta_H, \\ \theta_H, & \epsilon \geqslant \theta_H, \end{cases}$$

where $\hat{\epsilon}$ is the applied artificial viscosity, $\bar{h}$ is the arithmetic mean of $\mathbf{h}(\mathbf{x}) \in \mathbb{R}^d$, a local vector-measure of the element size described below, and $\mathcal{I} \in \mathbb{R}^m$ is a vector of ones. $\hat{\epsilon}$ is scaled to smoothly vary between zero and a maximum value, θ_H, as ϵ, the artificial viscosity produced by the shock capturing method, varies between θ_L, a minimum value, and θ_H. The determination of ϵ, based on a non-linear shock switch, will be described in greater detail in Section 4. θ_L and θ_H are defined to be proportional to $\lambda\bar{h}/p$, specifically, $\theta_L = 0.01\lambda_{\max}\bar{h}/p$ and $\theta_H = \lambda_{\max}\bar{h}/p$, and $\lambda_{\max}$ is the maximum wave speed of the system.

The local measure of element size is a linear variation throughout the computational mesh. Using continuous, linear, nodal basis functions, $\mathbf{h}(\mathbf{x})$ can be written as,

$$\mathbf{h}(\mathbf{x}) = \sum_{k=1}^{d+1}\mathbf{H}_k\varphi_k(\xi(\mathbf{x})),$$

where $\mathbf{H}_k \in \mathbb{R}^d$ is the average value of the bounding box vectors of all elements bordering the kth principal node of an element and φ_k is the nodal, linear basis function associated with the node. The arithmetic mean, $\bar{h}(\mathbf{x})$, is therefore a continuously varying scalar function throughout the domain,

$$\bar{h}(\mathbf{x}) = \frac{1}{d}\sum_{i=1}^{d}\sum_{k=1}^{d+1}H_{ki}\varphi_k(\xi(\mathbf{x})).$$

For compressible flow, the artificial viscosity matrix can be modified to better preserve the behavior of the shock transition given by the Euler equations. The Rankine–Hugoniot shock jump relations state that, for the steady Euler equations, total enthalpy is conserved across the shock [40]. Therefore, when dealing with the compressible Navier–Stokes equations, $\mathbf{A}_{ij}^\epsilon$ in (2) can be substituted with $\tilde{\mathbf{A}}_{ij}^\epsilon$, an artificial viscosity matrix designed to preserve total enthalpy.

Isenthalpic formulations of the Euler equations have long been considered in the computational community. Lytton [41] and Jameson [42] are two examples of numerical discretizations designed to preserve total enthalpy throughout the flow field. In the steady Euler equations, the energy and mass equations are identical if the constant factor, H, is removed from the energy equation. However, discrete schemes do not necessarily satisfy this property. A discrete solution with constant H is admissible if the numerical dissipation applied to the energy equation reduces to the numerical dissipation applied to the continuity equation when ρH is replaced by ρ [43]. Consequently, the application of the artificial viscosity matrix to the conservative state vector according to (2) would violate this criteria because the dissipation in the energy equation would act on ρE. One approach to derive an isenthalpic formulation of the artificial viscosity matrix is to apply the previous diagonal matrix, $\mathbf{A}_{ij}^\epsilon$ to a modified state vector which includes ρH instead of ρE. Specifically, $\tilde{\mathbf{A}}_{ij}^\epsilon$ is defined as,

$$\tilde{\mathbf{A}}_{ij}^\epsilon\nabla\mathbf{u} = \mathbf{A}_{ij}^\epsilon\nabla\tilde{\mathbf{u}}, \tag{3}$$

where $\tilde{\mathbf{u}} = [\rho, \rho v_i, \rho H]^T$.

Another formulation of an artificial viscosity matrix for the preservation of total enthalpy uses the Navier–Stokes viscosity matrix for an ideal gas. If the Prandtl number is set to, $Pr = 0.75$, one can show that this choice gives shock transitions with constant total enthalpy. This approach is used by Persson and Peraire [15], but is not applied to the results presented in this work.

In addition to the artificial viscosity matrix, numerical diffusion is added to the DG FEM scheme through the approximate flux function as well. For a flux function to ensure that total enthalpy is constant throughout the domain, it must also apply the same dissipation to the energy equation as the continuity equation multiplied by the constant factor, H. Jameson [43] describes and presents a few approximate flux functions that satisfy this criteria. In this work, two different approximate flux functions were employed. The first is the Roe flux function [44], which does not ensure constant total enthalpy, and the other is the modification of the van Leer flux difference splitting by Hänel et al. [45], which is designed to preserve total enthalpy.

The performance of both the isenthalpic artificial viscosity matrix, $\widetilde{\mathbf{A}}^{\epsilon}_{ij}$, and the non-isenthalpic diagonal matrix, $\mathbf{A}^{\epsilon}_{ij}$, is demonstrated in one of the numerical examples presented in Section 5. The same example also depicts the performance of the two approximate flux functions, the Roe flux function and the van Leer–Hänel flux function.

3. Motivation for smooth artificial viscosity

The one-dimensional viscous Burgers equation is employed to demonstrate the benefits of a smooth variation in artificial viscosity, compared to a non-smooth representation. The governing equation is modified to support a steady-state shock solution,

$$\frac{\partial}{\partial x}\left(\frac{1}{2}u^2\right) = \alpha u + \frac{\partial}{\partial x}\left(\nu(x)\frac{\partial u}{\partial x}\right) + f(x), \quad \text{for } x \in \Omega \subset \mathbb{R}, \tag{4}$$

where $u(x)$ is the state variable, $\nu(x)$ is the viscosity, α is a constant ($\alpha = -0.1$) and the forcing term, $f(x)$, is set such that the exact, steady-state solution with $\nu = 0$ has a shock at $x = 0$,

$$u(x) = \begin{cases} 2 + \sin\left(\frac{\pi x}{2}\right), & x < 0, \\ -2 - \sin\left(\frac{\pi x}{2}\right), & x > 0. \end{cases} \tag{5}$$

The viscosity, $\nu(x)$, is prescribed to be either a piecewise-constant or smooth Gaussian function, as depicted in Fig. 1. The piecewise-constant viscosity is applied to the cells immediately adjacent to the shock location with adjustable amplitude,

$$\nu(x) = \begin{cases} \beta \frac{h}{p} \|u(x)\|_\infty, & -h \leqslant x \leqslant h \\ 0, & \text{else}, \end{cases} \tag{6}$$

where β is an adjustable parameter. The Gaussian distribution of viscosity is specified to have a standard deviation equal to the cell size and the same total area as the piecewise-constant rectangle between $x \in [-h, h]$.

To perform the comparison, (4) is discretized using sixth order Legendre polynomials in a discontinuous Galerkin formulation. A global L^2-norm of the solution error and an H^1-norm of the error outside of the shock layer is measured between the discrete and exact solution for the two viscosity formulations. The error norms are defined as,

$$\|u - u_H\|_{L^2} = \left[\int_\Omega (u - u_H)^2\, dx\right]^{1/2},$$

$$\|u - u_H\|_{H^1} = \left[\int_{\Omega\setminus\delta_s} (u - u_H)^2 + \left(\frac{du}{dx} - \frac{du_H}{dx}\right)^2 dx\right]^{1/2},$$

where u_H is the discrete solution and δ_s is the shock layer, defined to be the distance extending from $x = 0$ to where the discrete solution is first within 0.5% of the exact solution.

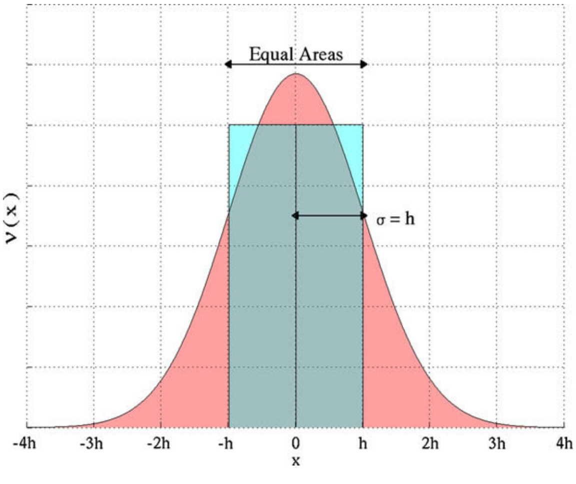

Fig. 1. Distributions of piecewise-constant and Gaussian artificial viscosity as applied to the 1D modified Burgers equation.

The comparison between the two viscosity representations is made at three different viscosity amplitudes ($\beta = 0.3,\ 1.0,\ 3.0$). The solution behavior for low, moderate, and high values of viscosity are shown in Fig. 2 with the error norm comparison in Table 1. At a low viscosity amplitude, the numerical oscillations in $u(x)$ are damped, but oscillations still remain in the derivative, $u_x(x)$, for both solutions, suggesting that the shock is not entirely resolved. At a higher viscosity amplitude, the Gaussian viscosity solution is smooth for both $u(x)$ and $u_x(x)$, but the piecewise-constant viscosity solution still has significant oscillations in $u_x(x)$. These oscillations are due to the conservation of the flux, $(u^2/2 + \nu u_x)$, across element boundaries. A jump in ν, requires a similar jump in u_x, but there is no change in u_x that can compensate for a jump to $\nu = 0$. Thus, for higher-order solutions, this jump in the viscosity induces fluctuations throughout the element. Finally, for much higher viscosity amplitudes, the Gaussian viscosity solution remains well-behaved, but the piecewise-constant solution suffers from oscillations in both $u(x)$ and $u_x(x)$.

Since the solutions for $u(x)$ are quite similar, except at the highest viscosity amplitude, the L^2-norm values of the solution error are also quite similar. However, the greater accuracy of the Gaussian viscosity solution is reflected in the H^1-norm values. At the high viscosity amplitude, the H^1-norm of the error for the Gaussian viscosity solution is smaller than the piecewise-constant solution by two orders of magnitude.

3.1. Multiple dimension issues

In one dimension, the errors induced by the shock for the non-smooth viscosity solution are generally confined to a region near the shock. In multiple dimensions, however, this is no longer the case. The jumps in viscosity from one element to the next along a discontinuity, due to changes in the shock strength, cell size and orientation, introduce errors both normal and tangential to the shock front. These errors create variations in the flow field which can convect downstream and pollute the

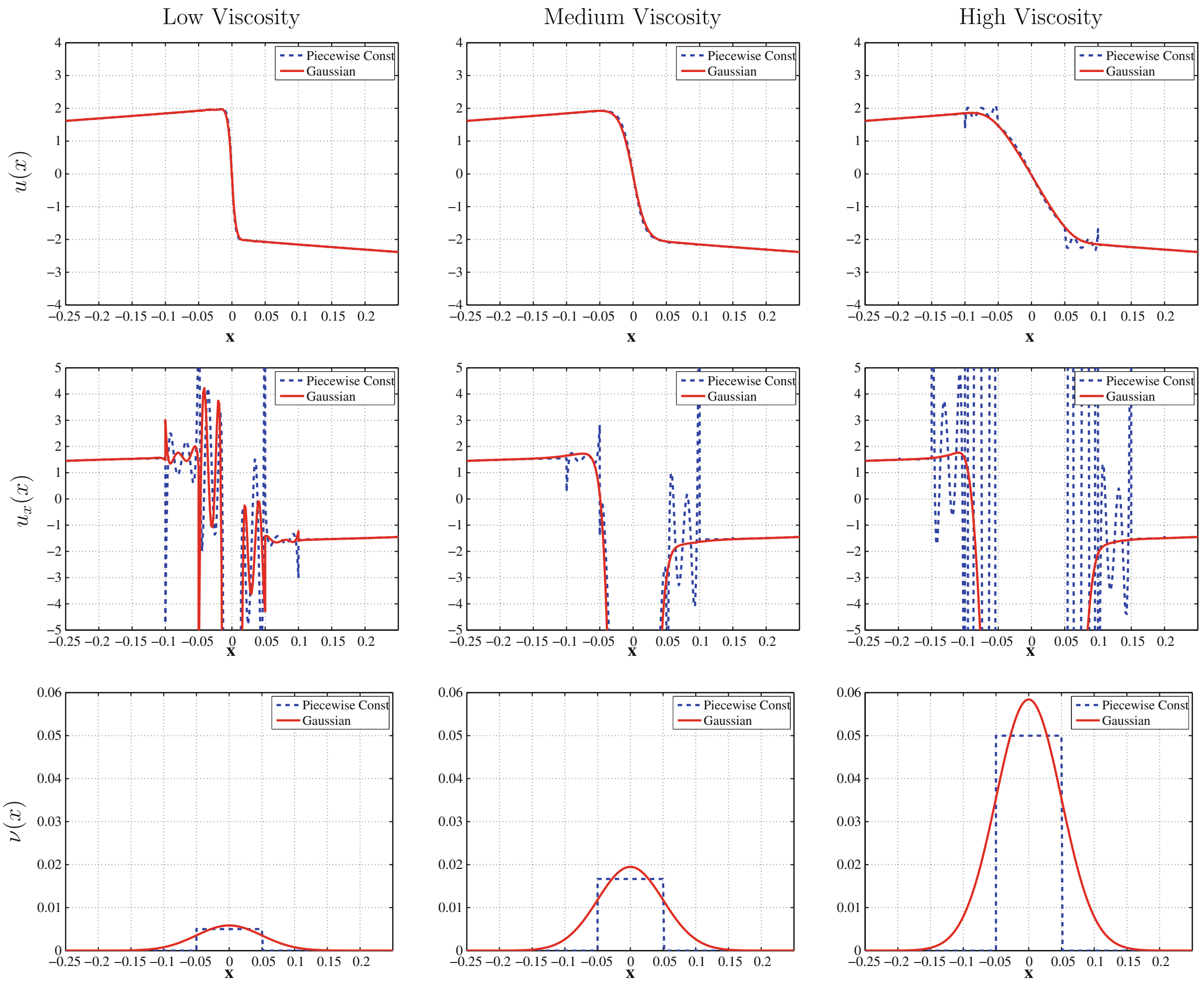

Fig. 2. Comparison of piecewise-constant and Gaussian viscosity solutions for modified Burgers equation across three different viscosity amplitudes (40 elements, $p = 6$).

Table 1
Comparison of L^2 and H^1 error norms for piecewise-constant and Gaussian viscosity solutions of modified Burgers equation (40 elements, $p = 6$).

	Low viscosity		Medium viscosity		High viscosity	
	H^1-norm	L^2-norm	H^1-norm	L^2-norm	H^1-norm	L^2-norm
Piece-const	0.873	0.124	0.680	0.227	19.196	0.387
Gaussian	0.548	0.134	0.180	0.244	0.167	0.405
β in (6)	0.3		1.0		3.0	

solution accuracy. For instance, consider the inviscid, supersonic flow over a cylinder depicted in Fig. 3. Using a elementwise-constant representation of artificial viscosity, significant oscillations in total pressure arise behind the shock front. While total pressure behind the bow shock is not constant, it should vary smoothly as the strength of the shock changes due to its curvature. The three lines shown in Fig. 3b are total pressure measurements taken from three grids across two uniform refinements along the solid black line in Fig. 3a. As the grid becomes finer, the total pressure oscillations persist. These variations in the solution downstream of a shock were previously observed by Quattrochi [46].

4. A PDE-based artificial viscosity model

In this section a PDE-based artificial viscosity model is proposed. We first describe the elementwise-constant (non-smooth) model. Then, extend this to the PDE-based formulation.

4.1. Non-smooth artificial viscosity

The non-smooth formulation of artificial viscosity is defined by an elementwise non-linear shock switch,

$$\epsilon|_\kappa = \frac{\bar{h}(\mathbf{x})}{p}\lambda_{\max}(\mathbf{u})S_\kappa(\mathbf{u}), \tag{7}$$

where $S_\kappa(\mathbf{u}) : \mathbb{R}^m \to \mathbb{R}$ is the non-linear switch or indicator function that detects the spurious numerical oscillations in element, κ, and determines the amount of artificial viscosity to add. The specifics of the shock switch are discussed in greater length in Section 4.3. As proposed by Persson and Peraire [15], the shock switch is multiplied by an h/p-scaling to allow for sub-cell shock resolution.

4.2. PDE-based artificial viscosity

To smooth the elementwise-constant artificial viscosity, a diffusion equation is chosen as the PDE model for artificial viscosity,

$$\frac{\partial \epsilon}{\partial t} = \nabla \cdot \left(\frac{\boldsymbol{\eta}}{\tau}\nabla\epsilon\right) + \frac{1}{\tau}\left[\frac{\bar{h}(\mathbf{x})}{p}\lambda_{\max}(\mathbf{u})S_\kappa(\mathbf{u}) - \epsilon\right] \quad \text{in } \Omega \subset \mathbb{R}^d \times \mathbb{R}^+, \tag{8}$$

$$\frac{\partial \epsilon}{\partial \hat{\mathbf{n}}} = \frac{\epsilon_\infty - \epsilon}{L} \quad \text{on } \partial\Omega, \tag{9}$$

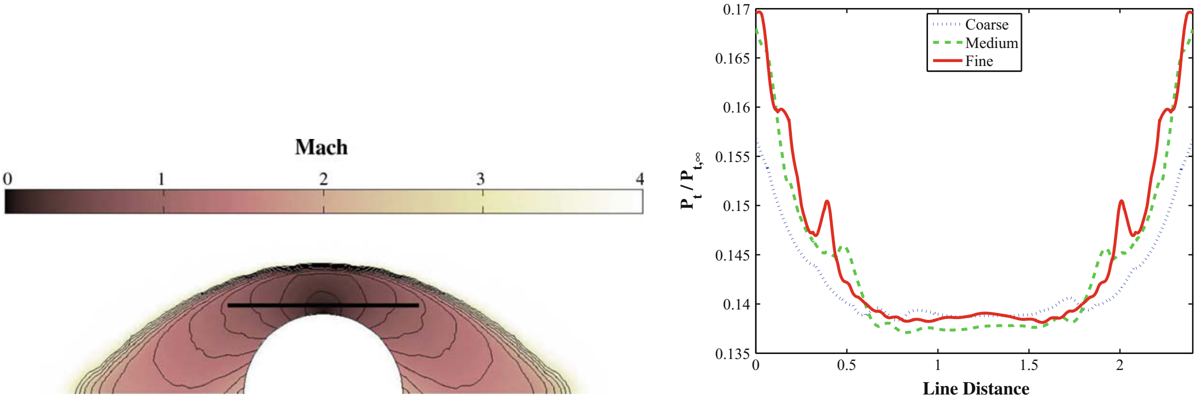

(a) Mach contours (b) Total pressure

Fig. 3. Mach number contours and total pressure measurements along a line behind the bow shock across two grid refinements for a $p = 3$ solution of a 2D flow around a cylinder at freestream Mach 4.

where τ is an appropriate time constant and $\boldsymbol{\eta} \in \mathbb{R}^{d\times d}$ is the diffusivity. The working variable of the PDE is $\epsilon(\mathbf{x}, t) : \mathbb{R}^d \times \mathbb{R}^+ \to \mathbb{R}$. For the sake of brevity, this PDE will be referred to as the *artificial viscosity equation* and its formulation of ϵ will be referred to as *PDE-based* artificial viscosity. The artificial viscosity is appended to the state vector of the Navier–Stokes equations and the artificial viscosity equation is discretized using the same DG approximation as described in Section 2.

The shock indicator, $S_\kappa(\mathbf{u})$, acts as a forcing term that drives ϵ to be non-zero in the vicinity of discontinuities. The diffusion term, governed by the parameter, $\boldsymbol{\eta}$, ensures that the viscosity is smooth (no large jumps at element edges) and that artificial viscosity produced in one element diffuses to its neighbors.

The artificial viscosity equation is cast with a time derivative and time constant, τ, defined such that ϵ evolves approximately as fast as the primary system of equations. This time scale is chosen to approximate the time it takes the fastest wave speed to traverse the resolution scale of the solution. To this end, the time constant, τ, is defined by

$$\tau = \frac{\min_i h_i}{C_1 p \lambda_{\max}(\mathbf{u})}; \quad C_1 = 3.$$

The second parameter of the artificial viscosity equation is the diffusion coefficient, $\boldsymbol{\eta} \in \mathbb{R}^{d\times d}$. For dimensional consistency it must have units of *Length*^{2}, and, so that the viscosity only spreads to neighboring elements, $\boldsymbol{\eta}$ should be made an explicit function of $\mathbf{h}(\mathbf{x})$. Thus, an appropriate setting of $\boldsymbol{\eta}$ is simply,

$$\boldsymbol{\eta} = C_2 \mathrm{diag}([h_x^2, h_y^2, h_z^2]^T),$$

The quantity, $\boldsymbol{\eta}/\tau$, is therefore,

$$\frac{\boldsymbol{\eta}}{\tau} = C_1 C_2 \frac{p\lambda_{\max}(\mathbf{u})}{\min_i h_i} \mathrm{diag}([h_x^2, h_y^2, h_z^2]^T); \quad C_1 C_2 = 15. \tag{10}$$

Since a shocks can occur at a boundary, and conceivably intersect a boundary at any angle, a boundary condition was sought that would not have significant impact on the behavior of the viscosity in the domain. Using a Green's function analysis, a Robin boundary condition was shown to work well, in which the gradient of ϵ is proportional to the difference between the boundary value and an ambient state ($\epsilon_\infty = 0$) over a local length scale, L ($L = 10\mathbf{h} \cdot \hat{\mathbf{n}}$) [39].

4.3. Shock indicators

The shock indicator, $S_\kappa(\mathbf{u})$, can take many forms. This research has employed two different indicators, which are presented in this section.

4.4. Resolution indicator

A resolution-based indicator was introduced by Persson and Peraire [15] as their method of detecting shocks to demonstrate the sub-cell shock capturing capabilities of artificial viscosity with higher-order, DG solutions. This indicator treats the higher-order solution as though it were comprised of a sequence of Fourier modes. For smooth flows, the coefficients of increasing Fourier modes should die away rapidly. In a true discontinuity, however, all frequency modes are present. This idea is similar to error indicators for adaptation in spectral methods [47]. With this concept in mind, a component of the state vector at any point in a higher-order approximation, can be represented as,

$$u(x) = \sum_{k=1}^{N(p)} U_k \phi_k(x),$$

where ϕ_k are the basis functions, U_k are the associated weights and $N(p)$ is the size of the higher-order expansion of degree p. Assuming an orthogonal basis,

$$\hat{u}(x) = \sum_{k=1}^{N(p-1)} U_k \phi_k(x),$$

where $\hat{u}$ is the truncated representation of the state vector at order $p - 1$.

With the definitions of u and $\hat{u}$, the resolution indicator can be defined by,

$$F_\kappa = \log_{10}\left(\frac{\left\langle f - \hat{f}, f - \hat{f} \right\rangle}{\langle f, f \rangle}\right), \tag{11}$$

where $\langle \cdot, \cdot \rangle$ represents the standard L^2 inner-product, and $f = f(\mathbf{u}) : \mathbb{R}^m \to \mathbb{R}$ is a component or function of the state vector. As with Persson and Peraire, this work relies on density as a reliable quantity for $f(\mathbf{u})$.

The final scaling of the indicator used by Persson and Peraire is such that it varies smoothly between zero and a maximum value,

$$S_\kappa(F_\kappa;\theta_S,\psi_0,\Delta\psi)=\begin{cases}0, & F_\kappa\leqslant\psi_0-\Delta\psi,\\ \theta_S, & F_\kappa\geqslant\psi_0+\Delta\psi,\\ \frac{\theta_S}{2}\left(1+\sin\frac{\pi(F_\kappa-\psi_0)}{2\Delta\psi}\right), & \|F_\kappa-\psi_0\|<\Delta\psi,\end{cases}$$

where θ_S is a maximum value ($\theta_S=1$) and ψ_0 and $\Delta\psi$ are empirically determined constants. In a Fourier expansion, coefficients decay at the rate of $1/p^2$ and since the resolution indicator works with the log of squared quantities, ψ_0 should roughly scale as $\psi_0\sim-4\log_{10}(p)$. In this work, $\psi_0\equiv-(4+4.25\log_{10}(p))$ and $\Delta\psi\equiv0.5$.

4.5. Jump indicator

The idea to use the uniquely DG inter-element jumps as a discontinuity indicator was first proposed by Dolejsi et al. [48] and also adopted by Krivodonova et al. [49] based on the work of Adjerid et al. [50]. For a smooth flow solution, the magnitude of the inter-element jumps should be convergent,

$$|[\![g]\!]|=\begin{cases}\mathcal{O}(h^{p+1}), & \text{smooth flow},\\ \mathcal{O}(1), & \text{discontinuity},\end{cases}$$

where $g=g(\mathbf{u}):\mathbb{R}^m\to\mathbb{R}$ is a state vector component or derived quantity,

Therefore, one can easily envision an indicator that measures jumps in a state quantity or a function to denote regions near a discontinuity. Specifically, the jump indicator is cast as,

$$J_\kappa=\frac{1}{|\partial\kappa|}\int_{\partial\kappa}\left|\frac{[\![g]\!]}{\{g\}}\right|\cdot\hat{\mathbf{n}}\,ds, \tag{12}$$

where jumps in pressure are chosen as the functional quantity, $g(\mathbf{u})$, to locate shocks. Additionally, for the purposes of linearization, the absolute value function was substituted with a $\mathcal{C}^1$-continuous approximation,

$$|x|\approx\frac{x^2}{\operatorname{sign}(x)x+\alpha},$$

where α is an input parameter. Similar to the resolution indicator, the final scaling of the jump indicator is smoothly limited by $S_\kappa=S_\kappa(J_\kappa;\theta_S,\psi_0,\Delta\psi)$, where θ_S is the same maximum value ($\theta_S=1$) and ψ_0 and $\Delta\psi$ are empirically determined constants, different from those of the resolution indicator. This work found that $\psi_0\equiv-(2.25+3\log_{10}(p))$ and $\Delta\psi\equiv0.5$ were reliable quantities.

It is important to note that the resolution indicator is a function of the state in a single element. The jump indicator, however, is dependent on the state values in neighboring elements as well. In the non-smooth viscosity approach to shock capturing, using the jump indicator with an otherwise compact discretization of diffusion terms would expand the numerical stencil of the entire scheme. This is because the artificial viscosity that is applied along element edges becomes dependent on the state values in immediate and second-degree neighbor elements as well. In contrast, with the artificial viscosity equation, the jump indicator is a source function and does not spread the numerical footprint of the scheme.

5. Numerical results

The above sections described two different artificial viscosity models: non-smooth and PDE-based. This section presents test cases designed to compare and contrast the performance of the two models.

5.1. Solution method

Although the focus of this work was on steady-state solutions, the unsteady term of the governing equations was retained to improve the initial transient behavior of the solver. Specifically, backward Euler time stepping was used with time step ramping. The restarted GMRES algorithm was used to iteratively solve the resulting linear system. To further aid iterative convergence of the linear system, an ILU factorization is used as a preconditioner where the factorization is performed using a reordering of elements into lines [51]. The lines are a unique set of elements created by the coupling between elements in a $p=0$ discretization of a scalar, linear convection–diffusion equation [38]. All of the higher-order solutions were arrived at via p-sequencing. Meaning, lower-order solutions served as the initial condition for higher-order solutions.

5.2. Convergence rate accuracy

5.2.1. Smooth flow

Both the resolution and jump indicators are designed to highlight under-resolved flow regions, such as those in the proximity of a discontinuity, that require the addition of artificial viscosity. However, for smooth, resolved flows, the shock indicators should not flag any troubled cells so that artificial viscosity is not unnecessarily added to the discretization.

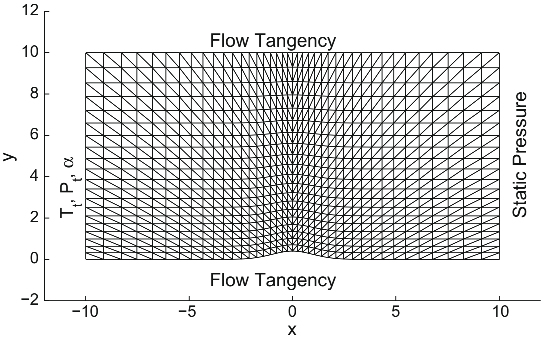

Fig. 4. Gaussian bump domain and mesh for smooth flow, shock indicator accuracy study (1600 elements).

The preservation of accuracy in smooth flow was tested on the problem of 2D, inviscid, subsonic flow over a Gaussian bump at a freestream Mach number of 0.5. An accuracy study of the entropy norm, $\|s - s_\infty\|_2$, was performed over five grids representing four uniform grid refinements from 400 to 102,400 elements. The 1600-element mesh is shown in Fig. 4. The total temperature, total pressure and flow angle were specified at the inflow boundary and the ambient static pressure was specified at the outflow boundary. Flow tangency was enforced at the upper and lower domain boundaries, and the bump surface was approximated using cubic geometry elements.

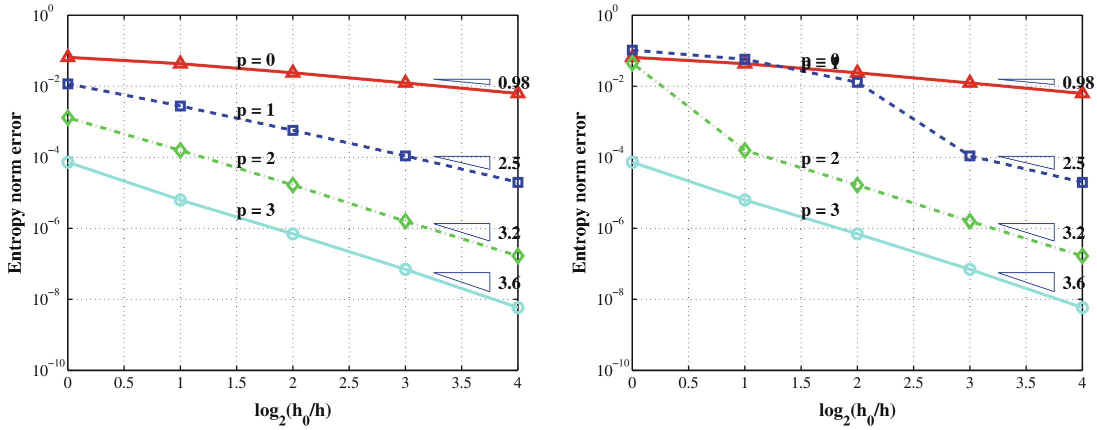

(a) No artificial viscosity (b) NS, Resolution indicator

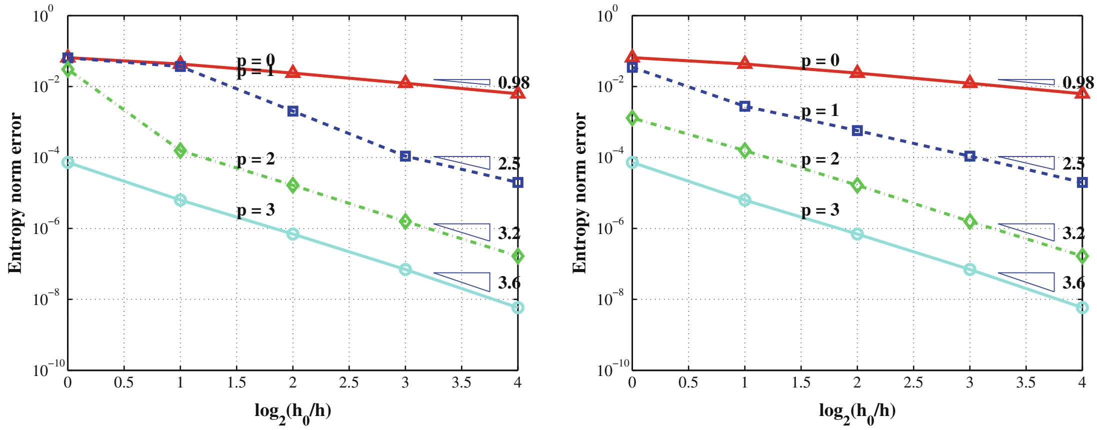

(c) PDE, Resolution indicator (d) PDE, Jump indicator

Fig. 5. Grid convergence rates of entropy norm for inviscid flow over a Gaussian bump, $M_\infty = 0.5$ with non-smooth (NS) and PDE-based artificial viscosity models.

The nominal grid convergence, without the use of shock capturing, is compared to the non-smooth and PDE-based artificial viscosity methods in Fig. 5 using both shock indicators. With sufficient flow field resolution, all of the artificial viscosity models achieve the same accuracy as the nominal results (without artificial viscosity), independent of the shock indicator used. For the low-order solutions on the coarser meshes, where the flow is not well resolved, the artificial viscosity added to the discretization increases the error relative to the nominal case.

5.2.2. Discontinuous flow

In addition to their smooth flow behavior, it is also desirable for the shock indicators and the artificial viscosity models to attain analytic convergence rates for flows with discontinuities as well. For discontinuous functions, the optimal, convergence rate in the L^1 norm for an optimal L^1 polynomial approximation is $\mathcal{O}(h/p)$ [52]. This is because the L^1-error is dominated by the discontinuity, which has zero thickness. Thus, the convergence rate in L^1 depends on how well the solution can approximate the thickness of the discontinuity, which is governed by the resolution length scale, h/p.

A verification of the analytic convergence rates for flows with discontinuities was carried out in the context of the modified Burgers equation with the discontinuous forcing function described in (4) and (5). However, instead of prescribing the artificial viscosity as an explicit function in space, the non-smooth and PDE-based artificial viscosity models were used. The results in Fig. 6 demonstrate that the optimal rate is achieved for shock flow cases. The grid convergence rates for the test case, across three uniform grid refinements, for $p = 1–4$ are $\mathcal{O}(h)$. Multiplying each line by their respective order, p, collapses all of the lines onto one, confirming the $\mathcal{O}(h/p)$ behavior. This is true for both the non-smooth and PDE-based artificial viscosity models and for both the resolution and jump indicators.

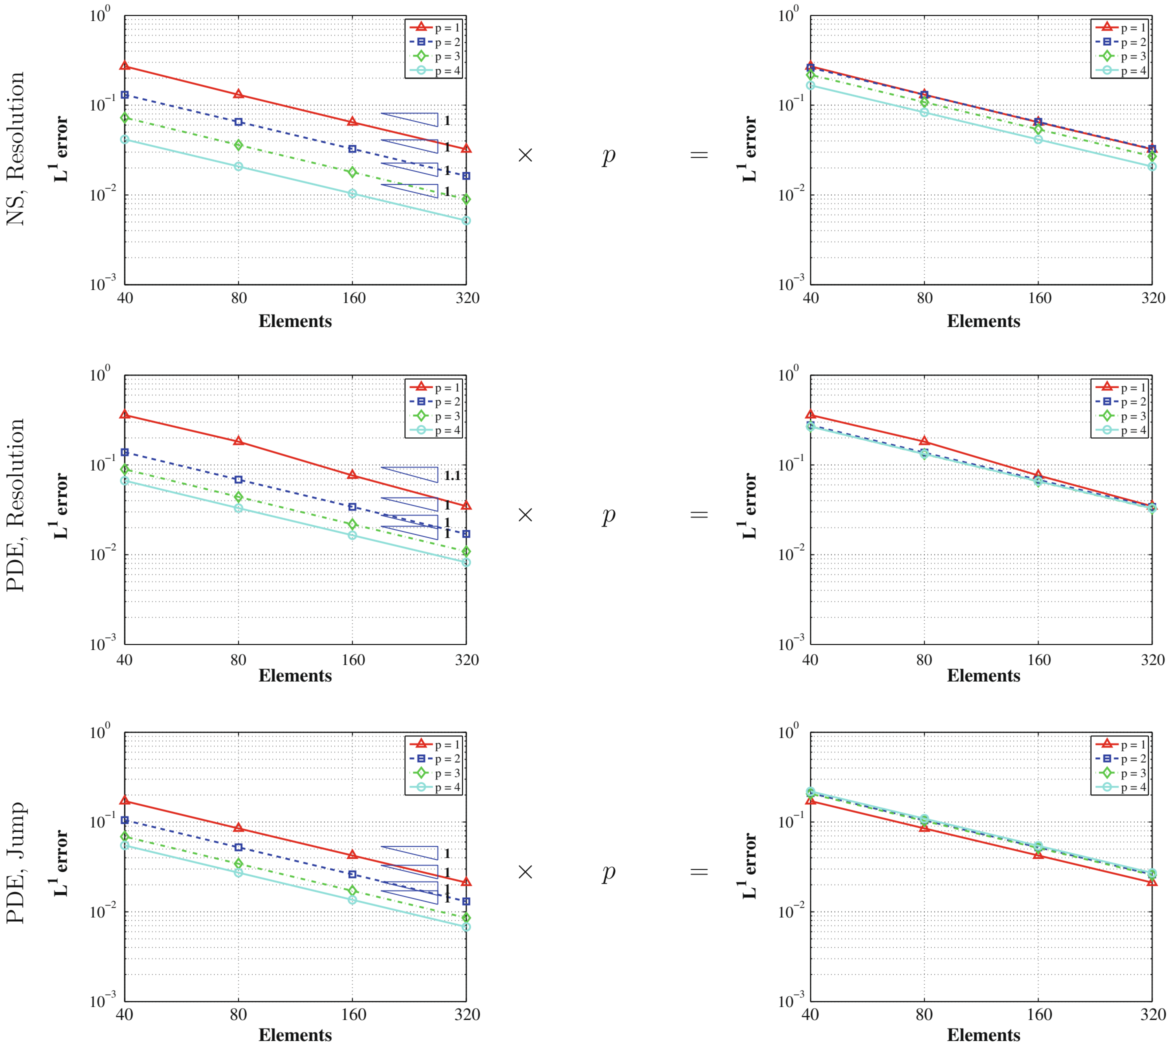

Fig. 6. L^1 grid convergence rates for 1D modified Burgers equation, (4) and (5), with both non-smooth (NS) and PDE-based artificial viscosity.

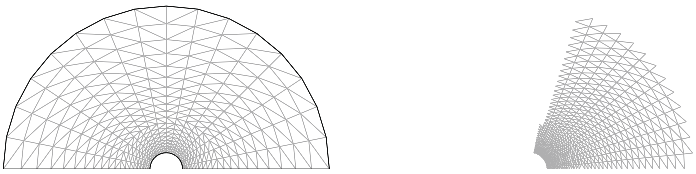

(a) Coarse

5.3. *Supersonic flow: half-cylinder, $M_\infty = 4$*

This test case is designed to examine the effectiveness of the artificial viscosity models for stronger shocks and, in particular, focus on the behavior of the stagnation quantities behind a shock front. The case also offers the opportunity to evaluate the efforts taken to conserve total enthalpy through a shock described in Section 2.2.

The solution for a 2D half-cylinder in a steady, inviscid, Mach 4 flow was solved on a sequence of three grids, shown in Fig. 7, representing two uniform refinements. The full state vector was specified at the inflow boundary and flow tangency was enforced on the cylinder surface. Since the outflow is supersonic, extrapolation boundary conditions were applied to the two outflow boundaries on either side of the cylinder.

The Mach number contours of the solution using the two artificial viscosity models and the resolution indicator are shown in the first row of Fig. 8. The remaining rows represent the variation of total enthalpy behind the shock along the solid black line in the Mach number plots. Using the Roe flux [44] and the Laplacian artificial viscosity matrix, $\mathbf{A}^{\epsilon}$ in (2), produces significant variations in total enthalpy behind the shock, with more oscillation in the non-smooth artificial viscosity than the PDE-based method. If the artificial viscosity matrix designed for the preservation of total enthalpy, $\tilde{\mathbf{A}}^{\epsilon}$ in (3), is applied instead, the variation in total enthalpy is significantly reduced for both viscosity models (notice the change in axis scaling). Finally, when the van Leer–Hänel flux function is used [45], the variation in total enthalpy is further damped. Changing the flux function has a smaller impact on the variation in total enthalpy because the relative amount of dissipation added to the scheme by the flux function is much less than the artificial viscosity matrix. The variations of total enthalpy behind the shock are more pronounced for the non-smooth artificial viscosity model than the PDE-based approach.

As described in Section 3.1, the total pressure is also impacted by the non-smooth viscosity model. Fig. 9 compares the variation of total pressure along the line behind the shock front for the two artificial viscosity models using the van Leer–Hänel flux function and the $\tilde{\mathbf{A}}^{\epsilon}$ viscosity matrix. The PDE-based artificial viscosity solution does not exhibit the oscillations that plague the non-smooth artificial viscosity solution. Instead, the total pressure varies smoothly as the shock strength changes due to curvature.

5.4. *Hypersonic flow: half-cylinder, $M_\infty = 17.605$, Re = 376,930*

The hypersonic flow over a half-cylinder is a simple test problem that highlights the difficulties in surface heat transfer prediction while using unstructured meshes in hypersonic flow. The problem is steady, features laminar flow everywhere in the boundary layer and the gas is modeled as thermally and calorically perfect. The complete description of the flow conditions are given in Table 2.

Gnoffo and White [4] first presented this half-cylinder test case and compared results obtained using the Langley Aerothermodynamic Upwind Relaxation Algorithm (LAURA) and the Fully Unstructured Navier–Stokes 3D (FUN3D) codes. LAURA is a structured mesh discretization designed for hypersonic, thermochemical non-equilibrium flows [53]. FUN3D is an unstructured finite volume method that includes the ability to perform error estimation, mesh adaptation, and design optimization for fluid dynamic problems [54].

Gnoffo and White computed flow over a two-dimensional half-cylinder extruded in the third dimension. For the LAURA computations, they created a structured hexahedral mesh which was adapted to align with the bow shock and included ten spanwise elements. The unstructured grid for FUN3D was generated directly from the structured mesh by uniformly biasing the diagonals of the cylinder surface mesh and then dividing the hexahedra into tetrahedra. The surface heat transfer distribution was sampled at ten different spanwise locations. Unlike the LAURA results, the FUN3D results were asymmetrical in both the circumferential and spanwise directions. Furthermore, the variations in heat transfer near the stagnation point were as high as 20%. Nompelis et al. [3] replicated this 3D test case and also found errors in the predicted heat transfer as high as 20%.

5.4.1. *2D structured grid results*

Four nested 2D structured meshes across three uniform refinements were used to conduct an *hp* grid convergence study. The coarsest mesh contained 4320 elements and the finest contained 276,480 elements. The domain and first mesh are shown in Fig. 10. Higher-order cubic geometry nodes were inserted for every element in the domain. The full state vector

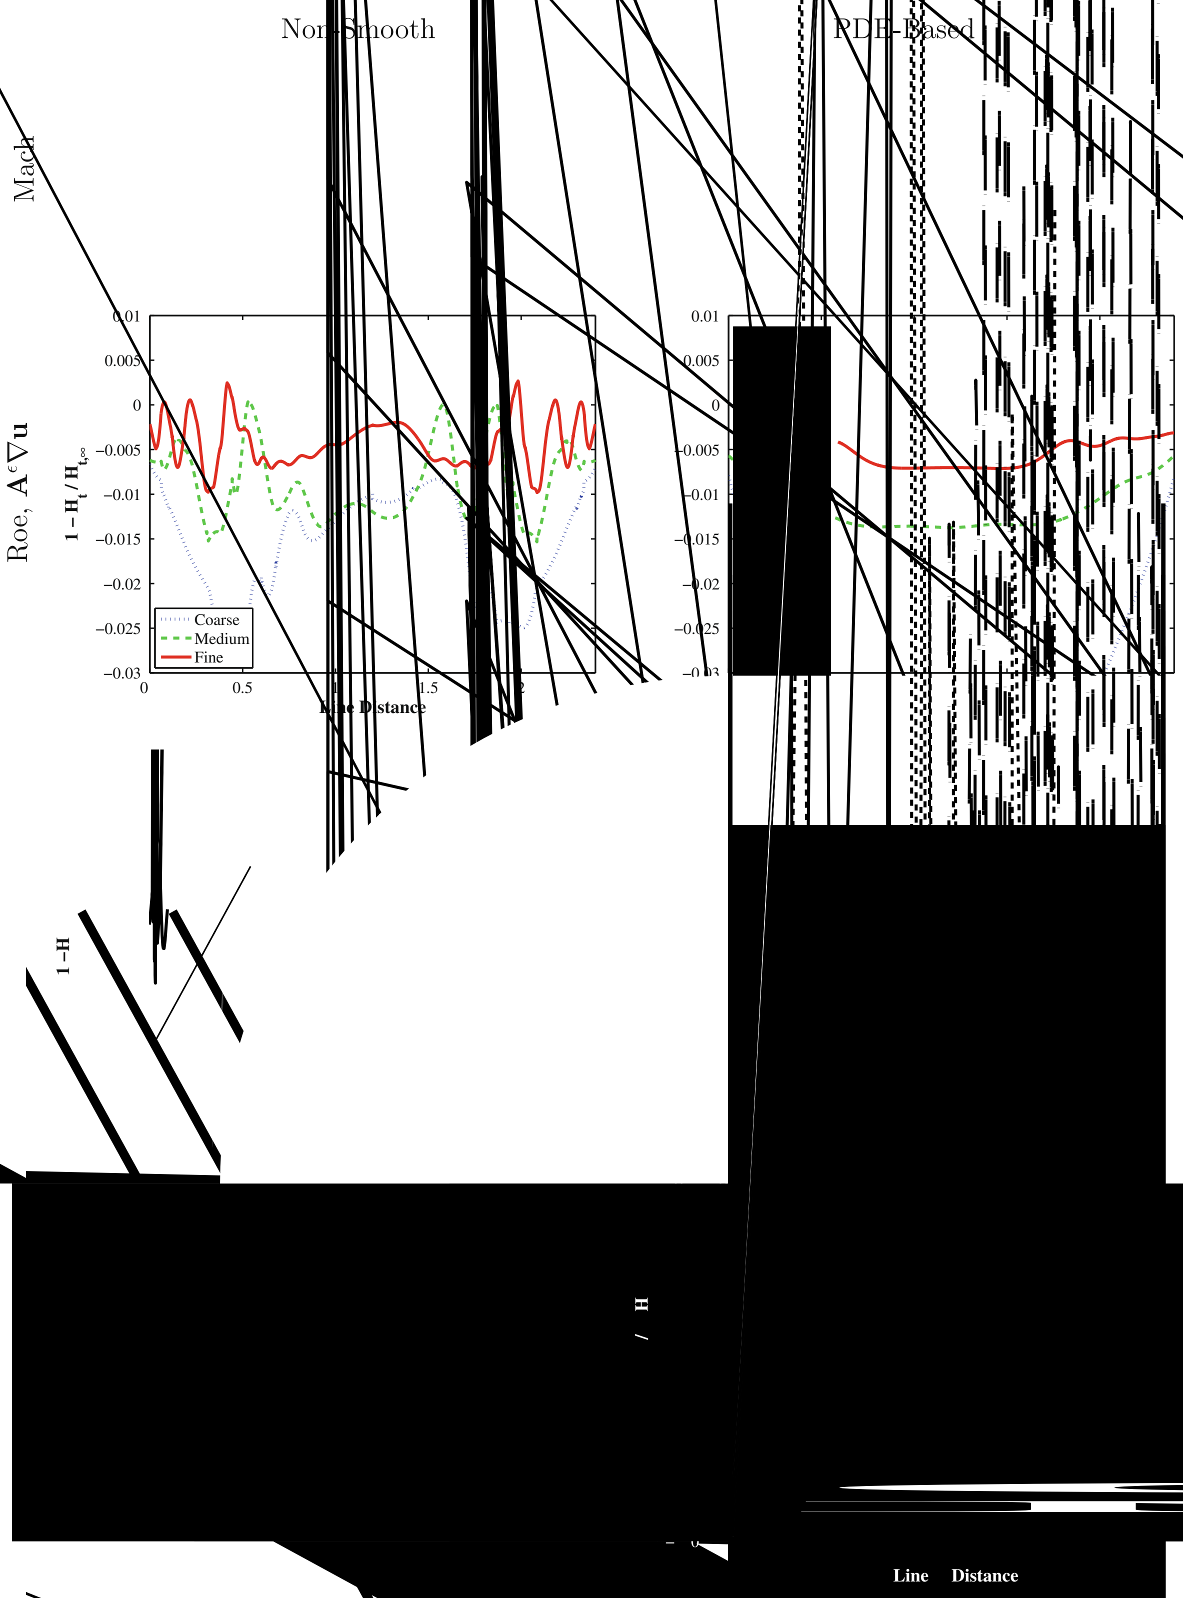

Fig. 8. Comparison of viscosity models with resolution

was specified at the inflow boundary and a no-slip condition at the constant wall temperature was enforced on the cylinder surface. The two outflow boundaries were extrapolation for supersonic flow.

The results of the grid convergence study are shown in Fig. 11. The plotted quantities are,

$$\text{Pressure coefficient}: \quad C_p = \frac{p - p_\infty}{\frac{1}{2}\rho V_\infty^2},$$

$$\text{Stanton number}: \quad C_h = \frac{q_{wall}}{c_p \rho_\infty V_\infty (T_{t,\infty} - T_{wall})},$$

where q_{wall} is the heat transfer at a surface, c_p is the constant pressure specific heat of the fluid, and $T_{t,\infty}$ is the freestream stagnation temperature.

The surface pressure distribution converges rapidly and agrees well with the LAURA results, even on the coarser solutions. The $p = 2$ and $p = 3$ solutions on the finest mesh agree well with the LAURA results for heat transfer. The coarser results tend to under-predict the peak heat transfer at the stagnation point due to the coarser resolution of the bow shock. The grid convergence study confirms that, with sufficient shock resolution, the higher-order DG solutions with PDE-based artificial viscosity can obtain accurate surface pressure and heating distributions on structured meshes. The grid converged, structured mesh results will serve as a baseline of comparison for the unstructured mesh results to follow.

5.4.2. *2D unstructured grid results*

For the 2D application, five hybrid grids were constructed, each with an identical, structured boundary layer mesh and with different unstructured meshes in the outer flow region. This focused the study on the influence of unstructured grid

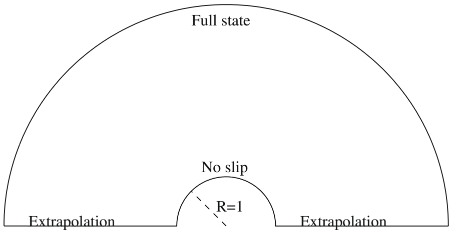

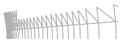

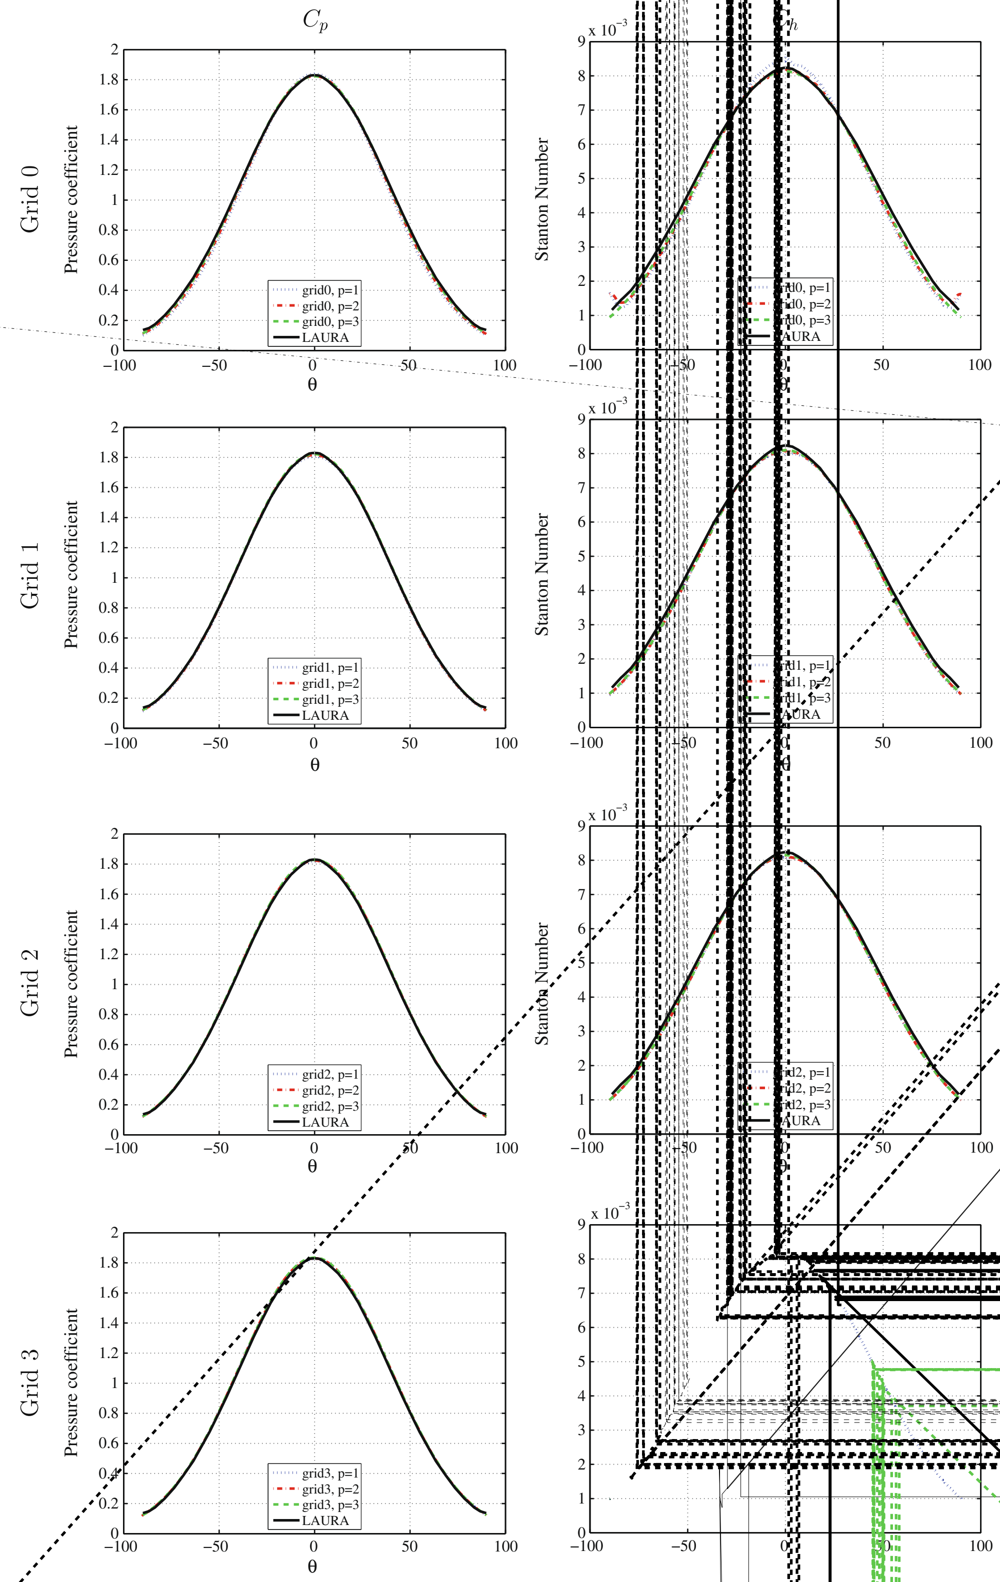

elements on shock resolution and on the downstream surface heat transfer. The structured boundary layer consisted of 61 nodes evenly spaced in the circumferential direction and 38 nodes in the radial direction and extended to 20% of the cylinder

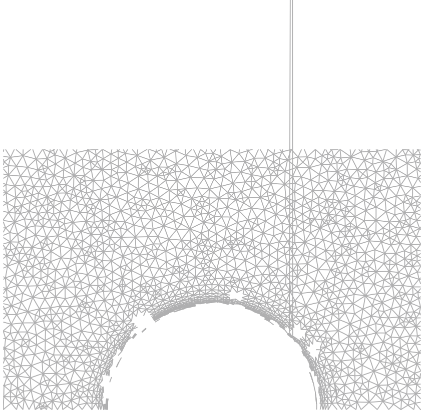

radius in the normal direction. To build the unstructured portion of the mesh, nodes were randomly placed in the external flow region until the maximum bounding box dimension was below a prescribed threshold (0.15 was used as the tolerance). To improve on mesh quality, after the addition of fifty randomly placed nodes, each node that was not located on a domain boundary was moved toward the center of mass of the polygon formed by the adjacent triangles. After the determination of the coordinates of the principal nodes, higher-order cubic geometry nodes were inserted for all of the elements. The near-field views of two of the five hybrid grids are shown in Fig. 12.

The temperature distribution of the flow field is shown in Fig. 13. Despite the coarse mesh in the vicinity of the shock and the misalignment with the bow shock trajectory, the shock is smoothly resolved. The significant temperature rise behind the bow shock can also be observed near the stagnation point.

Fig. 14 depicts surface quantities from the five different hybrid meshes and compares them to the LAURA results. In general, the DG surface quantities agree well with the LAURA results. For the surface pressure plot, there is excellent agreement

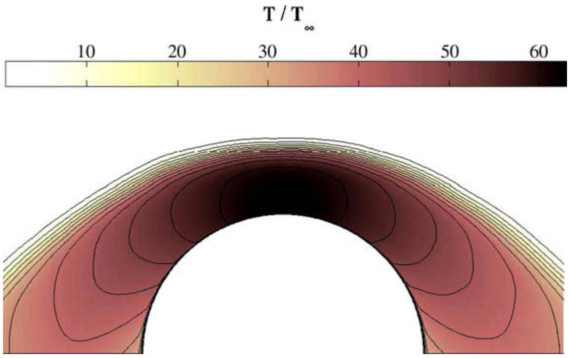

Fig. 13. Temperature distribution for $p = 3$ solution of flow over 2D half-cylinder with $M_\infty = 17.605$, Re = 376,930 (taken from Hybrid Grid 0).

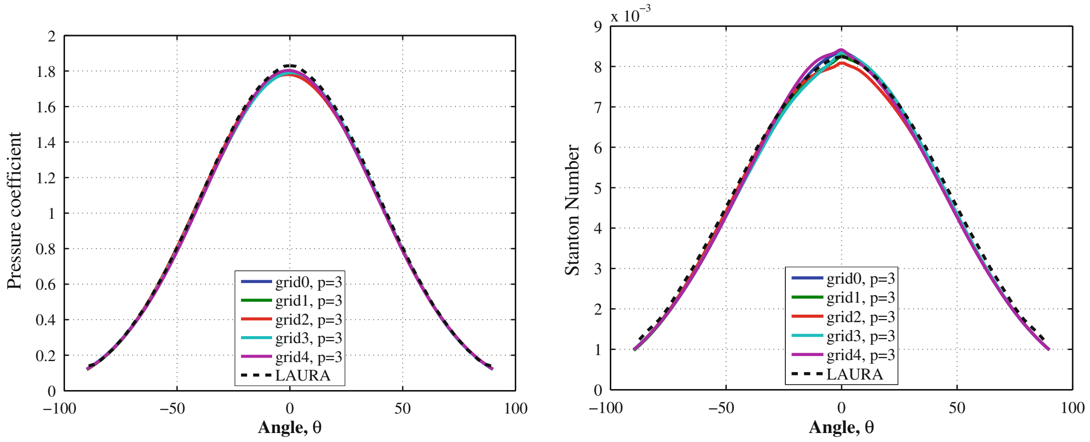

Fig. 14. Cylinder surface plots of $p = 3$ solution of flow over 2D half cylinder with $M_\infty = 17.605$, Re = 376,930 (all five grids shown).

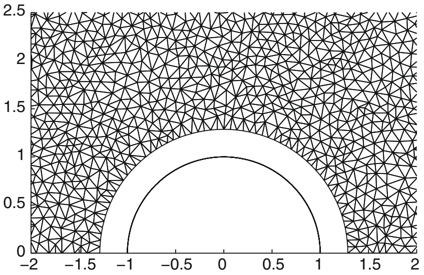

between the DG and LAURA results. The only discrepancy is the pressure at the stagnation point of the cylinder, where the DG results under-predict the LAURA results by about 2%. This is perhaps due to the coarser mesh in the vicinity of the shock for the DG results. For the Stanton number, there is a slight mismatch of the predicted peak heat transfer value and minor variation in the circumferential direction between the five different DG meshes. However, the variations are on the order of 1–2% and are far from the errors on the order of 20% reported by Gnoffo and White and Nompelis et al. This suggests that the higher-order solution coupled with the higher-order artificial viscosity distribution has significantly mitigated the sensitivity of the shock capturing to the grid cell size and orientation along the shock.

5.4.3. 3D unstructured grid results

The 3D mesh was also a hybrid grid, with a structured boundary layer mesh and an unstructured mesh in the external flow. The same structured node x–y coordinates as used in the 2D boundary layer mesh were used at six different, evenly spaced z-coordinate locations from $z \in [0, 0.5]$. This extruded the cylinder in the third dimension across five elements. The hexahedra created by the replicated node coordinates in the boundary layer were then split into tetrahedra. The remaining volume was filled using isotropic, unstructured tetrahedra created by TetGen [55]. Finally, higher-order cubic geometry nodes were inserted for each element in the grid. Perspectives of the 3D cylinder mesh are shown in Fig. 15. The boundary conditions are the same as the 2D problem in Fig. 12(a), with flow tangency applied to the two boundaries in the spanwise direction.

The temperature distribution is shown in Fig. 16. In this plot, half of the contours are taken from the plane $z = 0.0$ and the other are taken from the plane $z = 0.5$ to highlight any spanwise variation that might exist. As can be seen in the plots, the contours are well aligned and there is no significant spanwise variation.

The surface quantities of interest are plotted at six different spanwise locations in Fig. 17. The results are consistent with the 2D results above, and there is generally good agreement with the LAURA results. There is very little spanwise or circumferential variation in both the surface pressure and heat transfer distribution, with all spanwise locations agreeing quite well with the LAURA predictions, despite the fully unstructured mesh in the external flow.

6. Conclusion

This work presented a shock capturing methodology for higher-order ($p > 1$) finite element methods utilizing artificial viscosity. The drawbacks of a piecewise-constant artificial viscosity approach were highlighted. To achieve a smoother variation of artificial viscosity, this work proposed a PDE-based model of artificial viscosity which is solved in conjunction with the original governing equations. The method was implemented for the compressible Navier–Stokes equations using the discontinuous Galerkin finite element method. The algorithm was applied to hypersonic flow over a half-cylinder in both 2D and 3D. Using a series of hybrid structured–unstructured meshes (structured meshes for the boundary layer and unstructured meshes for the external flow), the higher-order solutions with the PDE-based artificial viscosity demonstrated good prediction of surface heating and were less susceptible to the errors introduced by the unstructured grid. In an upcoming paper, the approach is applied in an adjoint-based adaptive setting to a wide variety of shock-dominated flows.

Acknowledgments

The authors thank the Project X development team for the many contributions during the course of this work. This work was supported by the National Defense Science and Engineering Graduate Fellowship, the US Air Force Research Laboratory (USAF-3306-03-SC-0001), The Boeing Company, and the National Aeronautics and Space Administration (Co-op Agreement NNX07AC70A).

References

[1] S. Yoon, P.A. Gnoffo, J.A. White, J.L. Thomas, Computational challenges in hypersonic flow simulations, AIAA Paper 2007-4265, 2007.
[2] P.A. Gnoffo, Computational fluid dynamics technology for hypersonic applications, AIAA Paper 2003-2829, 2003.
[3] I. Nompelis, T.W. Drayna, G.V. Candler, Development of a hybrid unstructured implicit solver for the simulation of reacting flows over complex geometries, AIAA Paper 2004-2227, 2004.
[4] P.A. Gnoffo, J.A. White, Computational aerothermodynamic simulation issues for unstructured grids, AIAA Paper 2004-2371, 2004.
[5] P.A. Gnoffo, Simulation of stagnation region heating in hypersonic flow on tetrahedral grids, AIAA Paper 2007-3960, 2007.
[6] G. Candler, M. Barnhardt, T. Drayna, I. Nompelis, D. Peterson, P. Subbareddy, Unstructured grid approaches for accurate aeroheating simulations, AIAA Paper 2007-3959, 2007.
[7] J. von Neumann, R. Richtmyer, A method for the numerical calculation of hydrodynamic shocks, J. Appl. Phys. 21 (1950) 232–237.
[8] T.J.R. Hughes, L. Franca, M. Mallet, A new finite element formulation for computational fluid dynamics: I. Symmetric forms of the compressible Euler and Navier–Stokes equations and the second law of thermodynamics, Comput. Methods Appl. Mech. Eng. 54 (1986) 223–234.
[9] T.J.R. Hughes, M. Mallet, A. Mizukami, A new finite element formulation for computational fluid dynamics: II. Beyond SUPG, Comput. Methods Appl. Mech. Eng. 54 (1986) 341–355.
[10] T.J.R. Hughes, M. Mallet, A new finite element formulation for computational fluid dynamics: III. The generalized streamline operator for multidimensional advective–diffusive systems, Comput. Methods Appl. Mech. Eng. 58 (1986) 305–328.
[11] T.J.R. Hughes, M. Mallet, A new finite element formulation for computational fluid dynamics: IV. A discontinuity capturing operator for multidimensional advective–diffusive systems, Comput. Methods Appl. Mech. Eng. 58 (1986) 329–336.
[12] R. Hartmann, P. Houston, Adaptive discontinuous Galerkin finite element methods for the compressible Euler equations, J. Comput. Phys. 183 (2002) 508–532.
[13] R. Hartmann, Adaptive discontinuous Galerkin methods with shock-capturing for the compressible Navier–Stokes equations, Int. J. Numer. Methods Fluids 51 (2006) 1131–1156.
[14] S. Aliabadi, S.-Z. Tu, M. Watts, An alternative to limiter in discontinuous Galerkin finite element method for simulation of compressible flows, AIAA Paper 2004-76, 2004.
[15] P.-O. Persson, J. Peraire, Sub-cell shock capturing for discontinuous Galerkin methods, AIAA Paper 2006-112, 2006.
[16] W.H. Reed, T.R. Hill, Triangular mesh methods for the neutron transport equation, Technical Report LA-UR-73-479, Los Alamos Scientific Laboratory, 1973.
[17] B. Cockburn, G.E. Karniadakis, C.W. Shu, The development of discontinuous Galerkin methods, in: K. Cockburn, Shu (Eds.), Discontinuous Galerkin Methods: Theory, Computation and Applications, Springer, Berlin, 2000, pp. 3–50.
[18] T.C. Warbuton, I. Lomtev, Y. Du, S.J. Sherwin, G.E. Karniadakis, Galerkin and discontinuous Galerkin spectral/hp methods, Comput. Methods Appl. Mech. Eng. 175 (1999) 343–359.

[19] S.J. Sherwin, G.E. Karniadakis, A new triangular and tetrahedral basis for high-order finite element methods, Int. J. Numer. Methods Eng. 38 (1995) 3775–3802.
[20] B. Cockburn, C.W. Shu, Runge–Kutta discontinuous Galerkin methods for convection-dominated problems, J. Sci. Comput. 16 (3) (2001) 173–261.
[21] B. Cockburn, C.W. Shu, The local discontinuous Galerkin method for time-dependent convection–diffusion systems, SIAM J. Numer. Anal. 35 (6) (1998) 2440–2463.
[22] B. Cockburn, C.W. Shu, TVB Runge–Kutta local projection discontinuous Galerkin finite element method for scalar conservation laws: II. General framework, Math. Comput. 52 (1989) 411–435.
[23] B. Cockburn, S.Y. Lin, C.W. Shu, TVB Runge–Kutta local projection discontinuous Galerkin finite element method for conservation laws: III. One dimensional systems, J. Comput. Phys. 84 (1989) 90–113.
[24] B. Cockburn, S. Hou, C.W. Shu, TVB Runge–Kutta local projection discontinuous Galerkin finite element method for conservation laws: IV. The multidimensional case, Math. Comput. 54 (1990) 545–581.
[25] B. Cockburn, C.W. Shu, The Runge–Kutta discontinuous Galerkin finite element method for conservation laws: V. Multidimensional systems, J. Comput. Phys. 141 (1998) 199–224.
[26] S.R. Allmaras, A coupled Euler/Navier–Stokes algorithm for 2-D unsteady transonic shock/boundary-layer interaction, Ph.D. Thesis, Massachusetts Institute of Technology, 1989.
[27] S.R. Allmaras, M.B. Giles, A second-order flux split scheme for the unsteady 2-D Euler equations on arbitrary meshes, AIAA Paper 87-1119-CP, 1987.
[28] B. van Leer, Towards the ultimate conservative difference scheme. II: Monotonicity and conservation combined in a second order scheme, J. Comput. Phys. 14 (1974) 361–370.
[29] B. van Leer, Towards the ultimate conservative difference scheme. III: upstream-centered finite-difference schemes for ideal compressible flow, J. Comput. Phys. 23 (1977) 263–275.
[30] B. van Leer, Towards the ultimate conservative difference scheme: IV. A new approach to numerical convection, J. Comput. Phys. 23 (1977) 276–299.
[31] B. van Leer, Towards the ultimate conservative difference scheme: V. A second-order sequel to Godunov's method (for ideal compressible flow), J. Comput. Phys. 32 (1979) 101–136.
[32] F. Bassi, S. Rebay, High-order accurate discountinuous finite element solution of the 2-D Euler equations, J. Comput. Phys. 138 (1997) 251–285.
[33] F. Bassi, S. Rebay, A high-order accurate discountinuous finite element method for the numerical solution of the compressible Navier–Stokes equations, J. Comput. Phys. 131 (1997) 267–279.
[34] F. Bassi, S. Rebay, GMRES discontinuous Galerkin solution of the compressible Navier–Stokes equations, in: K. Cockburn, Shu (Eds.), Discontinuous Galerkin Methods: Theory, Computation and Applications, Springer, Berlin, 2000, pp. 197–208.
[35] K. Bey, J.T. Oden, A Runge–Kutta discontinuous finite element method for high speed flows, AIAA Paper 91-1575, June 1991.
[36] K. Bey, J.T. Oden, hp-version discontinuous Galerkin methods for hyperbolic conservation laws, Comput. Methods Appl. Mech. Eng. 133 (1996) 259–286.
[37] T.A. Oliver, Multigrid solution for high-order discontinuous Galerkin discretizations of the compressible Navier–Stokes equations, Master's Thesis, Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, August 2004.
[38] K.J. Fidkowski, T.A. Oliver, J. Lu, D.L. Darmofal, *p*-multigrid solution of high-order discontinuous Galerkin discretizations of the compressible Navier–Stokes equations, J. Comput. Phys. 207 (1) (2005) 92–113.
[39] G.E. Barter, Shock capturing with PDE-based artificial viscosity for an adaptive, higher-order, discontinuous Galerkin finite element method, Ph.D. Thesis, Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, June 2008.
[40] J. John, D. Anderson, Modern Compressible Flow, McGraw-Hill, Boston, MA, 1990.
[41] C.C. Lytton, Solution of the Euler equations for transonic flow over a lifting airfoil – the Bernoulli formulation (Roe/Lytton Method), J. Comput. Phys. 73 (1987) 395–431.
[42] A. Jameson, Solution of the Euler equations for two-dimensional transonic flow by a multigrid method, Appl. Math. Comput. 13 (1983) 327–356.
[43] A. Jameson, Analysis and design of numerical schemes for gas dynamics: 2. Artificial diffusion and discrete shock structure, Int. J. Comput. Fluid Dynam. 5 (1995) 1–38.
[44] P.L. Roe, Approximate Riemann solvers, parameter vectors, and difference schemes, J. Comput. Phys. 43 (1981) 357–372.
[45] D. Hänel, R. Schwane, G. Seider, On the accuracy of upwind schemes for the solution of the Navier–Stokes equations, AIAA Paper 87-1105, 1987.
[46] D.J. Quattrochi, Hypersonic heat transfer and anisotropic visualization with a higher-order discontinuous Galerkin finite element method, Masters Thesis, Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, 2006.
[47] C. Mavriplis, Adaptive mesh strategies for the spectral element method, Comput. Methods Appl. Mech. Eng. 116 (1994) 77–86.
[48] V. Dolejsi, M. Feistauer, C. Schwab, On some aspects of the discontinuous Galerkin finite element method for conservation laws, Math. Comput. Simul. 61 (3) (2003) 333–346.
[49] L. Krivodonova, J. Xin, J.-F. Remacle, N. Chevaugeon, J.E. Flaherty, Shock detection and limiting with discontinuous Galerkin methods for hyperbolic conservation laws, Appl. Numer. Math. 48 (3) (2004) 323–338.
[50] S. Adjerid, K.D. Devine, J.E. Flaherty, L. Krivodonova, A posteriori error estimation for discontinuous Galerkin solutions of hyperbolic problems, Comput. Methods Appl. Mech. Eng. 191 (2002) 1097–1112.
[51] L.T. Diosady, A linear multigrid preconditioner for the solution of the Navier–Stokes equations using a discontinuous Galerkin discretization, Masters Thesis, Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, May 2007.
[52] J.L. Guermond, A finite element technique for solving first-order PDEs in l^{p*}, SIAM J. Numer. Anal. 42 (2) (2004) 714–737.
[53] P.A. Gnoffo, F.M. Cheatwood, User's Manual for the Langley Aerothermodynamic Upwind Relaxation Algorithm (LAURA), NASA Langley Research Center, Virginia, NASA Technical Memorandum 4674, 1996.
[54] E. Nielsen, FUN2D/3D Fully Unstructured Navier–Stokes User Manual, NASA Langley Research Center, Computational Modeling and Simulation Branch, Virginia. <http://fun3d.larc.nasa.gov>.
[55] H. Si, TetGen, A Quality Tetrahedral Mesh Generator and Three-Dimensionsal Delaunay Triangulator, Weierstrass Institute for Applied Analysis and Stochastic, Berlin, Germany, 2006. <http://tetgen.berlios.de>.