



Shocks without shock capturing: Information geometric regularization of finite volume methods for Navier–Stokes-like problems

Anand Radhakrishnan^a, Benjamin Wilfong^a , Spencer H. Bryngelson^{a,b,c} ,¹, Florian Schäfer^d ,^{*},¹

^a School of Computational Science & Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA

^b Daniel Guggenheim School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA

^c George W. Woodruff School of Mechanical Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA

^d Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA

ARTICLE INFO

Dataset link: <https://github.com/sbryngelson/igr>, https://github.com/f-t-s/IGR_Manual_1dFVM

Keywords:

Computational fluid dynamics
compressible flow
Finite volume methods
Shock capturing
PDE regularization
Navier–Stokes equation

ABSTRACT

Shock waves in high-speed fluid dynamics produce near-discontinuities in the fluid momentum, density, and energy. Most contemporary works use artificial viscosity or limiters as numerical mitigation of the Gibbs–Runge oscillations that result from traditional numerics. These approaches face a delicate balance in achieving sufficiently regular solutions without dissipating fine-scale features, such as turbulence or acoustics. Recent work by Cao and Schäfer introduces information geometric regularization (IGR), the first inviscid regularization method for fluid dynamics. IGR replaces shock singularities with smooth profiles of adjustable width, without dissipating fine-scale features. This work provides a strategy for the practical use of IGR in finite-volume-based numerical methods. We illustrate its performance on canonical test problems and compare it against established approaches based on limiters and Riemann solvers. Results show that the finite volume IGR approach recovers the expected solutions in all cases. Across canonical benchmarks, IGR achieves accuracy competitive with WENO and LAD shock-capturing schemes in both smooth and discontinuous flow regimes. The IGR approach is computationally light, with meaningfully fewer memory accesses and arithmetic operations per time step.

1. Introduction

The problem: Shock waves in compressible CFD

Shock waves pose a central challenge in the numerical simulation of compressible fluids, particularly in high-speed fluid dynamics. Such waves introduce near-discontinuities in the momentum, density, and energy fields. These discontinuities cause Gibbs–Runge oscillations for numerical methods based on higher-order polynomial reconstruction [1–3]. Tailoring the mesh to the shock can overcome this problem, but it is prohibitively complicated for geometrically complex shock-interaction problems. Alternatively, the field of *shock capturing* aims to coarse-grain the effects of the microscopically thin shock without resolving it or interfering with the solution in other regions of the flow.

^{*} Corresponding author.

E-mail addresses: shb@gatech.edu (S.H. Bryngelson), florian.schaefer@nyu.edu (F. Schäfer).

¹ Equal contribution.

Shock capturing: Localized viscosity and limiters

Fluid dynamics studies relying on computation of compressible flow often add some form of localized artificial viscosity near the shock using so-called shock sensors [4–12]. This expands the shock to the grid scale, mitigating Gibbs–Runge oscillations. However, these viscous approaches must balance regularity against the preservation of fine-scale features, which often have physical meaning. Further, approaches based on local viscosity can break down at high Mach numbers [13].

Instead of regularizing the Navier–Stokes equations with viscosity, limiters lower the approximation order near shocks [1,14–19]. Limiters are often more robust than artificial diffusivity, but likewise need to balance the taming of shocks with the preservation of fine-scale features. Furthermore, they are intrusive changes to the numerical methods that can prohibit certain downstream applications. In particular, the nonlinear switching logic inherent to limiters can render adjoint-based sensitivity calculations undefined or introduce systematic biases [20,21], often requiring practitioners to accept errors due to “freezing” the limiters [22]. IGR’s entropic pressure, being the solution of a linear self-adjoint elliptic problem, does not share this drawback.

Information geometric regularization (IGR): The first inviscid shock regularization

Recent work by Cao and Schäfer [23] introduces the first regularization of compressible CFD that does not rely on viscosity. They first derive this *information geometric regularization* (IGR) in the pressureless inviscid (Euler) case, and observe that shock formation amounts to the collision and merging of particle trajectories. They then change the geometry of particle motion such that trajectories cannot cross, but instead asymptotically approach each other (see Fig. 2). This approach avoids the formation of grid-scale singularities while preserving post-shock behavior.

When forces due to pressure and viscosity are included, IGR yields a regularized conservation law that replaces singular shocks with grid-scale smooth profiles at a user-defined scale. The higher-order grid-scale smoothness of these profiles yields robust solutions even at high Mach numbers (see Wilfong et al. [24] for an application to a two-species Mach 14 flow). Because of its inviscid nature, this regularization has little impact on fine-scale features, such as turbulence or acoustic waves. In the unidimensional pressureless case, the entropic pressure equation reduces to $\Sigma/\rho - \alpha \partial_x (\Sigma/\rho) = 2\alpha (\partial_x u)^2$ and the regularized momentum equation reads $\partial_t(\rho u) + \partial_x(\rho u^2 + \Sigma) = 0$. This setting forms the starting point for the derivation of IGR by Cao and Schäfer [23]. In the unidimensional pressureless case, Cao and Schäfer [25] rigorously establish the existence of global smooth IGR solutions and their convergence to entropy solutions in the limit of vanishing regularization. Recent work of Barham et al. [26] investigates the connection between IGR and Hamilton regularizations as proposed by Guelmame et al. [27].

Outline

This work presents a finite-volume IGR implementation and illustrates its performance against canonical test problems. It is organized as follows. We begin in Section 2 by reviewing shock formation in the compressible Navier–Stokes equations and placing IGR in the context of viscous and limiter-based shock capturing. Section 3 introduces the IGR target equations and briefly sketches the geometric motivation. We highlight the practical bottlenecks addressed by IGR in Section 4, reviewing finite-volume methods and the roles of nonlinear reconstructions and Riemann solvers. Our finite-volume discretization of the IGR entropic pressure and the resulting algorithm are developed in Section 5. Numerical results are presented in Section 6. Section 7 closes with comparisons to related work, limitations, and conclusions.

2. Shock formation in the compressible Navier–Stokes equations

2.1. Compressible Navier–Stokes equations

The compressible Navier–Stokes equations are

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0, \quad (1)$$

$$\frac{\partial(\rho \mathbf{u})}{\partial t} + \nabla \cdot (\rho \mathbf{u} \mathbf{u}^T + p \mathbf{I} - \boldsymbol{\tau}) = 0, \quad (2)$$

$$\frac{\partial(E)}{\partial t} + \nabla \cdot [(E + p)\mathbf{u} - \boldsymbol{\tau} \cdot \mathbf{u}] = 0, \quad (3)$$

which represent the conservation of mass, momentum, and energy. Their variables include pressure p , density ρ , total energy E , time t , velocity $\mathbf{u} = \{u_1, u_2, u_3\}$ (in the 3D case); $\mathbf{I}$ is the identity tensor. To simplify the presentation, we use the ideal gas law equation of state

$$p = (\gamma - 1)\rho e, \quad \text{for } e = E/\rho - \|\mathbf{u}\|^2/2, \quad (4)$$

where γ is the ratio of specific heats (adiabatic exponent). However, the method is not restricted to this choice. The coefficients μ and ζ are the shear and bulk viscosity, and the viscous stress tensor is

$$\tau_{ij} = \mu \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) + (\zeta - 2\mu/3) \delta_{ij} \frac{\partial u_k}{\partial x_k}, \quad (5)$$

where δ_{ij} is the Kronecker delta, x_i are the coordinate directions and $i, j \in \{1, 2, 3\}$ for the 3D case.




Fig. 1. Comparison of localized artificial diffusivity (LAD) and information geometric regularization (IGR) approaches for shock treatment (modified from [24], with author permission.)

2.2. Shock formation

The nonlinear terms of the compressible Navier–Stokes equations cause wavefronts to steepen; faster-moving regions overtake slower ones, producing shocks (near-discontinuities) in pressure, density, and velocity. Viscosity smooths these features, but physical shocks remain thin (typically of the order of the mean free path [28]), far below typical Navier–Stokes grid resolution scales (and, indeed, the range of validity of the equations themselves). Thus, in the case of the Navier–Stokes equations, some degree of additional treatment via modeling or numerics is required to represent shock behavior faithfully. These additions modify the governing equations or the numerical methods to coarse-grain the shock to practical grid scales, a procedure known as shock capturing.

2.3. Viscous PDE regularizations

The most basic approach to regularizing the Navier–Stokes equations for shock handling is to add small viscous terms [4], modifying Eqs. (1)–(3) as

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = \varepsilon_\rho \nabla^2 \rho \quad (6)$$

$$\frac{\partial (\rho \mathbf{u})}{\partial t} + \nabla \cdot (\rho \mathbf{u} \mathbf{u} + p \mathbf{I} - \boldsymbol{\tau}) = \varepsilon_{\rho u} \nabla^2 (\rho \mathbf{u}) \quad (7)$$

$$\frac{\partial (E)}{\partial t} + \nabla \cdot [(E + p) \mathbf{u} - \boldsymbol{\tau} \cdot \mathbf{u}] = \varepsilon_E \nabla^2 (E), \quad (8)$$

where ε is the artificial viscosity coefficient. The value of ε determines the trade-off between shock capturing and the preservation of other flow features, like acoustics. More sophisticated approaches make this trade-off more palatable by letting ε depend on the current-time flow solution. The resulting nonlinear viscosity acts as a *shock sensor*, localizing it to the vicinity of the shock.

A notable viscous approach is localized artificial diffusivity (LAD) [6,8], which regularizes the Navier–Stokes equations as

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0 \quad (9)$$

$$\frac{\partial (\rho \mathbf{u})}{\partial t} + \nabla \cdot (\rho \mathbf{u} \mathbf{u}^T + p \mathbf{I} - \boldsymbol{\tau}) = \nabla \cdot (\zeta^* \nabla \cdot \mathbf{u}) \quad (10)$$

$$\frac{\partial (E)}{\partial t} + \nabla \cdot [(E + p) \mathbf{u} - \boldsymbol{\tau} \cdot \mathbf{u}] = \nabla \cdot (\zeta^* \nabla \cdot \mathbf{u}) \cdot \mathbf{u}. \quad (11)$$

Rather than adding a viscous regularization to each component of the conservation law, it modifies the equation by locally increasing its bulk viscosity. It is modulated by a function of the local flow state that detects under-resolved features (such as shocks), a so-called shock sensor. Using artificial bulk (as opposed to shear) viscosity ζ^* helps restrict additional dissipation to shocks rather than to the vortical fine-scale structures typical of turbulence. Many variants exist, but a representative example proposed by Mani et al. [8] is to choose $\zeta^* \propto \rho \max(0, -\nabla \cdot \mathbf{u})$, where the overline indicates the application of Gaussian smoothing with length scale commensurate to the grid spacing. The constant of proportionality is taken to be of the order of the squared grid spacing.

2.4. A discrete alternative: Limiters

Rather than regularizing the PDE, an alternative approach changes the numerical discretization to handle discontinuities. ENO schemes [29] adaptively select the smoothest stencil at each cell interface. WENO [1,16] generalizes this by forming a convex combination of candidate stencils, weighted by smoothness indicators. MUSCL [30,31] reconstructs piecewise-linear profiles within each cell. All three families achieve high-order accuracy in smooth regions while suppressing spurious oscillations near shocks.

These and related Total Variation Diminishing (TVD) schemes employ limiters, that is, nonlinear functions (for example, minmod, superbee, van Leer, MC, or the stencil selection in (W)ENO) that reduce the reconstruction order near discontinuities. The choice of limiter governs the trade-off between numerical dissipation and shock sharpness: more diffusive limiters, like minmod, are robust but smear shocks, and less diffusive ones, like superbee, preserve sharper shocks but are less stable. A promising alternative is convex limiting [18,19], which applies the minimum correction necessary to preserve invariant domain properties of the PDE.




Fig. 2. Information geometric regularization (IGR). In the pressureless case, IGR regularizes particle trajectories $t \mapsto \Phi_t(x)$ such that instead of colliding, particles asymptotically approach each other ((a), modified from [23], with author permission). With pressure, this replaces the singular shocks with smooth profiles on a scale $\sqrt{\alpha}$ (b). In the $\alpha \rightarrow 0^+$ limit, one recovers the Euler equations, and the auxiliary IGR equation Eq. (12) has the trivial solution $\Sigma \equiv 0$.

3. Information geometric regularization

3.1. The IGR target equations

The information geometric regularization of Eqs. (1)–(3) complements the physical pressure p with the *entropic pressure* Σ [23], defined by the auxiliary elliptic problem

$$\rho^{-1} \Sigma - \alpha \operatorname{div}(\rho^{-1} \nabla \Sigma) = \alpha (\operatorname{tr}^2([\nabla \mathbf{u}]) + \operatorname{tr}([\nabla \mathbf{u}]^2)), \quad (12)$$

where $\nabla \mathbf{u}$ is the flow Jacobian matrix with components $(\nabla \mathbf{u})_{ij} := \partial_j u_i$ and $\alpha > 0$ is the regularization parameter controlling the width of the smoothed shock profile. This yields the information geometric regularization

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0, \quad (13)$$

$$\frac{\partial(\rho \mathbf{u})}{\partial t} + \nabla \cdot (\rho \mathbf{u} \mathbf{u}^T + (p + \Sigma) \mathbf{I} - \boldsymbol{\tau}) = \mathbf{0}, \quad (14)$$

$$\frac{\partial(E)}{\partial t} + \nabla \cdot [(E + (p + \Sigma)) \mathbf{u} - \boldsymbol{\tau} \cdot \mathbf{u}] = 0. \quad (15)$$

3.2. Inviscid regularization with smooth solutions

Eqs. (13)–(15) have solutions that replace shocks with profiles that are smooth at the scale of $\sqrt{\alpha}$. At the same time, these profiles propagate and interact with other flow structures and with each other as if they were nominal shocks regularized by vanishing viscosity. The absence of viscous regularization means that fine-scale flow features, such as acoustic or entropy waves and turbulence features, are preserved.

IGR does not act as a diffusive mechanism in the conventional sense. Rather than damping the solution over time, it adds a pressure-like term that prevents the trajectories of gas volumes from colliding. The elliptic source $\operatorname{tr}^2([\nabla \mathbf{u}]) + \operatorname{tr}([\nabla \mathbf{u}]^2)$ is largest wherever velocity gradients are sharpest, so Σ grows in those regions and adds to the pressure that opposes the trajectory crossing, as illustrated in Fig. 2. As described in Cao and Schäfer [23, Section 7.1], IGR can be interpreted as a “nonlocal and sign-indefinite viscosity”. The nonlocality of Σ enables IGR to slow down gas particles approaching the shock *before they reach it*, thus enabling IGR to produce the higher-order smooth profiles shown in Fig. 1. The sign-indefinite nature of IGR ensures that it only activates in systematically compressive flows (shocks), but not in oscillatory features due to sound waves or turbulence, as illustrated in Fig. 1. This is the distinction from viscous regularizations, which damp all flow gradients regardless of their character.

3.3. A primer on the geometric derivation of IGR

We denote as Φ_t the *flow map*, which tracks where each fluid parcel, initially at position $\mathbf{x}$, moves to at time t . The flow map Φ_t satisfies $\mathbf{u}(\Phi_t(\mathbf{x})) = \dot{\Phi}_t(x)$ and $\rho(\Phi_t(x)) = \det[D\Phi_t(x)]$, where $D\Phi_t$ denotes the deformation gradient of the flow map, distinct from the velocity gradient $\nabla \mathbf{u}$ in Eq. (12). Solutions to Eqs. (1) and (2) describe a time evolution on the manifold of possible flow maps [32,33].

In the pressureless and inviscid case ($p = 0, \boldsymbol{\tau} = \mathbf{0}$), $t \mapsto \Phi_t$ describes a straight line on the manifold of diffeomorphisms (differentiably invertible maps). Shocks form when Φ_t reaches the boundary of this manifold, meaning that two points $\mathbf{x}_1 \neq \mathbf{x}_2$ are mapped to the same $\Phi_t(\mathbf{x}_1) = \Phi_t(\mathbf{x}_2)$. After shock formation, the trajectories $t \mapsto \Phi_t(\mathbf{x}_1)$ and $t \mapsto \Phi_t(\mathbf{x}_2)$ merge; Φ_t evolves on the

boundary of the diffeomorphism manifold. Cao and Schäfer [23] regularize this equation by changing the notion of “straight lines” on the diffeomorphism manifold such that they instead asymptotically approach the boundary, but do not reach it.

With IGR, trajectories $t \mapsto \Phi_t(\mathbf{x})$ merge asymptotically, without ever colliding, as shown in Fig. 2. This avoids the formation of singularities while preserving the long-time behavior. When adding back the pressure to the resulting equations, their solutions replace singular shocks with smooth profiles on a scale determined by the parameter α that measures the strength of the regularization.

3.4. IGR in practice

As outlined by Cao and Schäfer [23], a computational promise of IGR is obviating the need for nonlinear numerical methods, like limiters or Riemann solvers, for representing shock behavior. IGR produces no regularization in pure shear: the flow Jacobian there is traceless and nilpotent, so both $\text{tr}^2(\nabla \mathbf{u})$ and $\text{tr}(\nabla \mathbf{u}^2)$ on the right-hand side of Eq. (12) vanish and $\Sigma = 0$. IGR thus targets compressive features without affecting shear-dominated structures. IGR is formulated to reduce the numerical artifacts encountered by these approaches, such as artificial dissipation or carbuncles. At the same time, replacing nonlinear limiters and stencil-selection logic with standard linear reconstruction and flux evaluation reduces the computational cost of simulating flows with shocks, as demonstrated by Wilfong et al. [24]. Note that the full IGR scheme is nonlinear: the nonlinearity enters through the entropic pressure Σ , whose elliptic right-hand side depends quadratically on the velocity gradient. This work presents IGR in the same finite-volume framework used by traditional approaches to compressible flows. We next review this framework and identify the bottlenecks that IGR resolves.

4. Finite volume methods

Finite volume methods (FVMs) discretize the computational domain into control volumes and store cell-averaged solution values. Conservation is enforced by balancing fluxes across cell faces via the divergence theorem, making FVMs well-suited for problems involving shocks and other discontinuous fields.

4.1. Fundamental concept

In a finite volume formulation, the conservation law $\partial_t \varphi + \nabla \cdot \mathbf{F} = S$ is integrated over each control volume Ω . Applying the divergence theorem converts the volume integral of the divergence into a surface integral of outward-pointing face fluxes, yielding the semi-discrete balance

$$\frac{d}{dt} \int_{\Omega} \varphi \, d\Omega = S_{\Omega} + \sum_{f \in \{\text{faces of } \Omega\}} F_f A_f, \quad (16)$$

where F_f is the average flux through the face f , and A_f is that face's area. High-order reconstruction reduces the mismatch between left- and right-hand interface states, decreasing the dissipation introduced by the Riemann solver. Some Riemann solvers (for example, Lax–Friedrichs) entail more dissipation, encouraging numerical stability, while others (for example, HLLC, discussed next) are more aggressive.

4.2. Discrete finite volume formulation

We describe a typical finite-volume solver of the one-dimensional compressible Euler equations. The same strategy extends directly to the 3D compressible Navier–Stokes equations on a Cartesian grid by successively computing fluxes in each coordinate direction. One stores each conserved variable as a cell average, evaluates numerical fluxes at every cell face, and applies the 1D update formula along each coordinate direction.

The 1D compressible Euler equations can be expressed as a conservation law,

$$\frac{\partial \mathbf{U}(x, t)}{\partial t} + \frac{\partial}{\partial x} \mathbf{F}(\mathbf{U}(x, t)) = \mathbf{0}, \quad (17)$$

with conservative state variables $\mathbf{U}$ and fluxes $\mathbf{F}(\mathbf{U})$ as

$$\mathbf{U} = \begin{pmatrix} \rho \\ \rho u \\ E \end{pmatrix} \quad \text{and} \quad \mathbf{F}(\mathbf{U}) = \begin{pmatrix} \rho u \\ \rho u^2 + p \\ (E + p)u \end{pmatrix}, \quad (18)$$

$I_i = [x_{i-1/2}, x_{i+1/2}]$ is a grid cell of width Δx and the cell average at time t^n is

$$U_i^n = \frac{1}{\Delta x} \int_{I_i} \mathbf{U}(x, t^n) dx. \quad (19)$$

Evaluating Eq. (19) and time integrating Eq. (17) requires the fluxes at the finite volume faces $\mathbf{F}(\mathbf{U}_{L, R})$. The pressure depends on the state variables through the equation of state. Here, we consider the case of polytropic gas laws for simplicity, with an equation of state of the form $p = (\gamma - 1)(E - \rho u^2/2)$ with adiabatic exponent γ .

The standard first-order scheme for computing the fluxes at the half-cell distances uses piecewise constant reconstruction, with resulting left and right interface states $U_{i+1/2}^L = U_i$ and $U_{i+1/2}^R = U_{i+1}$. To achieve higher-order accuracy, we replace these constants by values obtained from a polynomial reconstruction of degree 5 constructed from neighboring cell averages and then evaluate the polynomial at the cell faces. For example, a fifth-order accurate polynomial reconstruction uses the reconstruction weights

$$w^L = \{2/60, -13/60, 47/60, 27/60, -3/60\}, \quad (20)$$

$$w^R = \{-3/60, 27/60, 47/60, -13/60, 2/60\}, \quad (21)$$

and defines the reconstructions as

$$U_{i+1/2}^L = \sum_{j=-2}^{+2} w_j^L U_{i+j}, \quad U_{i+1/2}^R = \sum_{j=-1}^{+3} w_j^R U_{i+j}. \quad (22)$$

The numerical method's spatial order of accuracy matches the accuracy of this polynomial reconstruction.

The numerical fluxes are denoted $\hat{F}_{i+1/2}$, and their evaluation depends on the numerical method used, including the Riemann solver. For this work, we only fully express the local Lax–Friedrichs (LF) Riemann solver, which is also known as the Rusanov numerical flux. With $U_L = U_{i+1/2}^L$ and $U_R = U_{i+1/2}^R$, the LF numerical flux is

$$\hat{F}_{i+1/2} = \frac{1}{2} [F(U_L) + F(U_R)] - \frac{\lambda_{i+1/2}}{2} (U_R - U_L), \quad (23)$$

where

$$\lambda_{i+1/2} = \max\{|u_L| + c_L, |u_R| + c_R\} \quad (24)$$

is the largest local wave speed, u is the velocity, and c is the sound speed. Integrating Eq. (17) over I_i and the time interval t^n to t^{n+1} gives the update equation for the state variables:

$$U_i^{n+1} = U_i^n - \frac{\Delta t}{\Delta x} (\hat{F}_{i+1/2}^n - \hat{F}_{i-1/2}^n). \quad (25)$$

This numerical scheme is stable so long as the Courant–Friedrichs–Lewy (CFL) condition is satisfied:

$$0 < \frac{\Delta t}{\Delta x} \max_i \lambda_{i+1/2} \leq 1. \quad (26)$$

The diffusive term in Eq. (23) is designed to recover a method sufficiently stable for transport-dominated problems. In smooth flow fields, the difference $(U_R - U_L)$ vanishes as the fifth power of the mesh width, limiting the spurious dissipation in this setting.

In low-regularity situations, such as shock formation, high-order reconstructions cause oscillations due to Gibbs–Runge phenomena. Treating discontinuous behavior has heretofore required limiters such as ENO [29] and MUSCL [30,31] schemes that locally reduce the approximation order, or other viscous mechanisms. The resulting lower order reconstructions increase the difference $(U_R - U_L)$ and thus the diffusive effects of Lax–Friedrichs Riemann solvers, as illustrated in Fig. 3.

Mitigating these effects requires the use of sophisticated approximate Riemann solvers to compute fluxes at cell interfaces, such as Harten–Lax–van Leer (HLL) and HLLC [34–36]. Those also involve nonlinear computations and intermediate variables, such as wave speeds. Together, the nonlinear computation of interface reconstructions and Riemann solvers requires that large temporary arrays be stored in device memory. Furthermore, they prevent the fusion of compute kernels that could otherwise compute these quantities in a single kernel pass. The confluence of these features prevents the aggregation of compute kernels that update grid-level quantities (such as fluxes) in place, resulting in computation and memory overhead. In contrast, the discrete flow Jacobian required by IGR uses a compact 3-point stencil and adds no temporary storage; the dominant savings come from eliminating the nonlinear stencil selection, smoothness indicators, and characteristic decompositions of WENO, not from avoiding gradient computation per se. Quantitative performance comparisons are provided in Wilfong et al. [24]. Beyond these computational issues, the reliance on grid-aligned treatments used by these Riemann solver-driven approaches is suspected to give rise to the carbuncle phenomenon [37]. Linear schemes, such as the Lax–Friedrichs method discussed in Section 4.2, avoid this shortcoming.

IGR replaces shocks with smooth solutions [25]. As illustrated in Fig. 3, this enables the use of higher-order reconstructions throughout, and thus foregoes nonlinear numerics in the form of limiters or Riemann solvers.

5. Solving the IGR equations with finite volume methods

Empirical and theoretical evidence suggest that the information geometrically regularized Eqs. (13)–(15) produce solutions that are smooth on the scale $\sqrt{\alpha}$ (Fig. 2, [23,25]). Thus, higher-order polynomial reconstructions avoid the formation of Gibbs–Runge oscillations for a mesh width $\lesssim \sqrt{\alpha}$. As a result, no limiters are needed; the difference between left and right reconstructions vanishes to high order in the mesh width, and the Lax–Friedrichs scheme is only weakly diffusive.




Fig. 3. Comparison of high-order polynomial reconstruction strategies in FVM near a jump discontinuity (orange). Blue horizontal segments are cell averages; the orange arrow marks the target face. (a) Fifth-order reconstruction from smooth data yields a single consistent polynomial. (b) The same stencil near a jump produces oscillatory left/right reconstructions. (c) (W)ENO selects among candidate sub-stencils (dotted), preferring the smoothest one. (d) IGR smooths the cell averages themselves, enabling a single high-order reconstruction without nonlinear stencil selection.

5.1. Discretization of the entropic pressure

The only change needed for these benefits is the addition of the entropic pressure Σ to the physical pressure p . Since the regularization term is only of order $\alpha \approx \Delta x^2$, we find that a piecewise constant approximation of Σ yields satisfactory results in practice, which are shown later in Section 6. Denoting ρ_{ijk} as the average density and $\rho \mathbf{u}_{ijk}$ as the average momentum over the ijk -th cell, and letting Δ_x , Δ_y , and Δ_z denote the grid spacings in the x , y , and z coordinate directions, we define the discrete flow Jacobian $\nabla \mathbf{u}_{ijk}$ at the ijk -th cell using the central stencil

$$\nabla \mathbf{u}_{ijk} = \left(\frac{1}{2\Delta_x} \left(\frac{\rho \mathbf{u}_{(i+1)jk}}{\rho_{(i+1)jk}} - \frac{\rho \mathbf{u}_{(i-1)jk}}{\rho_{(i-1)jk}} \right), \frac{1}{2\Delta_y} \left(\frac{\rho \mathbf{u}_{i(j+1)k}}{\rho_{i(j+1)k}} - \frac{\rho \mathbf{u}_{i(j-1)k}}{\rho_{i(j-1)k}} \right), \frac{1}{2\Delta_z} \left(\frac{\rho \mathbf{u}_{ij(k+1)}}{\rho_{ij(k+1)}} - \frac{\rho \mathbf{u}_{ij(k-1)}}{\rho_{ij(k-1)}} \right) \right), \quad (27)$$

which is, again, a linear and grid-local operation.

Solving the elliptic equation Eq. (12) for Σ using the finite difference discretization of the Laplacian yields

$$\begin{aligned} \frac{\Sigma_{ijk}}{\rho_{ijk}} + \alpha \left(\sum_{|i-\tilde{i}|=1} \frac{2(\Sigma_{ijk} - \Sigma_{\tilde{i}jk})}{\Delta_x^2 (\rho_{ijk} + \rho_{\tilde{i}jk})} + \sum_{|j-\tilde{j}|=1} \frac{2(\Sigma_{ijk} - \Sigma_{i\tilde{j}k})}{\Delta_y^2 (\rho_{ijk} + \rho_{i\tilde{j}k})} + \sum_{|k-\tilde{k}|=1} \frac{2(\Sigma_{ijk} - \Sigma_{ijk\tilde{k}})}{\Delta_z^2 (\rho_{ijk} + \rho_{ijk\tilde{k}})} \right) \\ = \alpha \left(\text{tr}^2(\nabla \mathbf{u}_{ijk}) + \text{tr}(\nabla \mathbf{u}_{ijk}^2) \right) \end{aligned} \quad (28)$$

Numerous methods exist for solving this system of equations. One choice is Jacobi iteration, which amounts to computing the next iterate Σ_{ijk}^{n+1} from the previous one Σ_{ijk}^n using the formula

$$\Sigma_{ijk}^{(n+1)} = \frac{\alpha \left(\text{tr}^2(\nabla \mathbf{u}_{ijk}) + \text{tr}(\nabla \mathbf{u}_{ijk}^2) \right) + \alpha \left(\sum_{|i-\tilde{i}|=1} \frac{2\Sigma_{ijk}^{(n)}/\Delta_x^2}{\rho_{ijk} + \rho_{\tilde{i}jk}} + \sum_{|j-\tilde{j}|=1} \frac{2\Sigma_{ijk}^{(n)}/\Delta_y^2}{\rho_{ijk} + \rho_{i\tilde{j}k}} + \sum_{|k-\tilde{k}|=1} \frac{2\Sigma_{ijk}^{(n)}/\Delta_z^2}{\rho_{ijk} + \rho_{ijk\tilde{k}}} \right)}{\frac{1}{\rho_{ijk}} + \alpha \left(\sum_{|i-\tilde{i}|=1} \frac{2/\Delta_x^2}{\rho_{ijk} + \rho_{\tilde{i}jk}} + \sum_{|j-\tilde{j}|=1} \frac{2/\Delta_y^2}{\rho_{ijk} + \rho_{i\tilde{j}k}} + \sum_{|k-\tilde{k}|=1} \frac{2/\Delta_z^2}{\rho_{ijk} + \rho_{ijk\tilde{k}}} \right)}. \quad (29)$$

Jacobi iteration is an effective iterative method for solving the elliptic equation Eq. (28), and is shown to be sufficient to compete with established numerical methods in Section 6. Other iterative methods, such as Gauss–Seidel or successive over-relaxation (SOR), can further accelerate this step. In our experience, the elliptic solve is not the computational bottleneck. After the first time step, the iteration is warm-started from the previous time step's Σ , so very few iterations are required to converge. Combined with the small value of $\alpha \propto \Delta x^2$ and the absence of global communication, the solve is well-suited to GPU execution. Quantitative cost comparisons with WENO are provided in Wilfong et al. [24].

5.2. Boundary conditions for Σ

The elliptic problem that defines Σ requires boundary conditions. We use Neumann boundary conditions on Σ , enforcing $\partial_n \Sigma = 0$ at the wall. This is the boundary-normal component of $\text{div}(\Sigma \mathbf{I})$, so Σ contributes no spurious normal force at the wall. Equivalently, the wall acts as a mirror plane for Σ , which is the correct condition for a shock propagating parallel to a flat wall. Algorithmically, this restricts the sums of Eqs. (28) and (29) to $1 \leq \tilde{i} \leq m_x, 1 \leq \tilde{j} \leq m_y, 1 \leq \tilde{k} \leq m_z$, where m_x, m_y , and m_z are the number of grid cells in each coordinate direction. The exception to this choice is rotation-symmetric or periodic boundary conditions. There, the elliptic problem uses the same wrap-around boundary conditions as the overall problem.

5.3. Smoothing boundary and initial conditions

Shock treatment via IGR is based on the principle that solutions of the regularized equations are as grid-scale smooth as their initial conditions. For such grid-scale smooth initial conditions, one recovers grid-scale smooth solutions always, mitigating Gibbs–Runge oscillations in higher-order reconstructions. However, many canonical problems in fluid dynamics, such as shock tubes, have discontinuous initial or boundary conditions. Without limiters or viscous regularizations, discontinuities of initial or boundary conditions will trigger Gibbs–Runge oscillations that pollute the simulation.

This problem can be addressed by applying a small amount of smoothing to ensure smoothness *on the grid scale* of the initial and boundary conditions. In practice, one can apply a small, constant number of diffusive pseudo-time steps to the state variables U in the initial and boundary conditions to achieve this. IGR preserves the smoothness of the initial and boundary conditions at the grid scale. Thus, different from the filtering steps in spectral methods, *no further smoothing needs to be applied during the simulation*.

5.4. Discretization of the resulting equation

The above-described steps ensure that the PDE solutions remain smooth, enabling the use of higher-order accurate reconstructions, such as the fifth-order accurate ones in Eq. (20). The high-order accuracy of the reconstructions allows using the straightforward Lax–Friedrichs Riemann solver while incurring little numerical diffusion. Because IGR solutions are smooth at the scale $\sqrt{\alpha}$, the left/right interface mismatch satisfies $\|U_R - U_L\| = O((\Delta x/\sqrt{\alpha})^5)$. Thus, the LF dissipation term is also $O((\Delta x/\sqrt{\alpha})^5)$ and can be controlled by an appropriate choice of $\sqrt{\alpha}$. Numerous time integrators are available for the resulting semi-discretization. Strong-stability-preserving Runge–Kutta schemes perform well in practice [38]; we use them for the numerical experiments in Section 6.

6. Numerical results

6.1. Convergence

Using PDE-based regularizations, such as IGR, raises three distinct questions about the convergence of the resulting numerical method. The first question is whether the numerical solution converges as $\Delta x \rightarrow 0$ for a fixed α (and thus, a fixed PDE) to the continuous solution of the regularized PDE. The second is whether the exact solution of the regularized PDE converges as the regularization parameter tends to zero ($\alpha \rightarrow 0$). Third, and most practically relevant, is the convergence of the numerical approximation of IGR as both the regularization parameter and the grid scale tend to zero, following their natural scaling relationship ($\sqrt{\alpha} \propto \Delta x \rightarrow 0$).

Fig. 4 demonstrates convergence in these three regimes on a one-dimensional test case with initial condition $u(x) = 1.5 \sin(2\pi x)$, $\rho(x) \equiv 1, e(x) = 4$, distinguishing behavior before and after shock formation. For a fixed α , we observe second-order convergence before shock formation and almost second-order convergence after shock formation. Since we are using a second-order accurate discretization of the entropic pressure Σ , this matches the expected order of accuracy. When decreasing α for a fixed and sufficiently small Δx , we observe second-order (in $\sqrt{\alpha}$) convergence before shock formation and first-order convergence after a shock has formed. The former is to be expected because the entropic pressure is of the order $\alpha(\partial_x u)^2$. The first-order convergence in the shock arises because in a shock of width $\sqrt{\alpha}$, the flow derivative $\partial_x u$ is of the order $\alpha^{-1/2}$, and thus the entropic pressure is constant to leading order. Thus, the overall decrease of the L^1 norm of the entropic pressure is only due to the decreasing width of the shock, resulting in first-order convergence. Similar behavior is observed in virtually all approaches to shock treatment in Eulerian fluid simulations, except for shock tracking [39]. When letting $\sqrt{\alpha} \propto \Delta x \rightarrow 0$ jointly, we again observe second-order convergence before shock formation and first-order afterward. In light of the results above, the observed convergence rate is consistent with the theoretical convergence rate.




Fig. 4. We investigate the convergence of IGR and its numerical solution in three different regimes. In (a), we fix α and observe the convergence with decreasing grid size. We see (near) second-order convergence both before and after the shock forms. In (b), we observe the convergence of IGR solutions with decreasing α , computed on a fine grid. We observe second-order (in $\sqrt{\alpha}$) convergence before shock formation, and first-order convergence thereafter. In (c), we investigate the behavior as both α and Δx tend to zero. We choose the scaling relationship $\sqrt{\alpha} \propto \Delta x$, which results in shock widths proportional to the grid size, as would be expected in a practical application of our approach. We again observe second-order convergence before and first-order convergence after shock formation.

6.2. Example cases

We consider representative 1D–3D test problems that highlight IGR’s behavior in the presence of shocks and fine-scale wave content. For 1D smooth-wave propagation, we consider a sinusoidal wave benchmark. For shock-dominated flows, we consider the Shu–Osher interaction problem and shock-tube benchmarks. In 2D, we consider a Riemann interaction with added entropy-wave content and the Double Mach reflection problem. In 3D, we consider compressible Taylor–Green turbulence and a multi-jet high-Mach configuration. For baseline comparisons, we use fifth-order WENO reconstruction in two variants. Characteristic WENO projects the conserved variables into the local characteristic fields of the flux Jacobian before applying the WENO stencil selection. We combine it with a Roe Riemann solver [36]. Component WENO applies WENO independently to each conserved variable. We combine it with a Rusanov (local Lax–Friedrichs) Riemann solver. The characteristic decomposition typically reduces spurious oscillations near discontinuities, and the Roe approximate Riemann solver is less diffusive than the Rusanov one, but they entail higher computational costs. Wilfong et al. [24] report a fourfold speedup using IGR and replacing WENO with linear reconstruction of the same order and HLLC with Rusanov.

In many applications, it is most convenient to smooth the initial conditions (see Section 5.3) via diffusion. For the sake of having a well-defined and grid-independent initial condition, we instead replace jumps by scaled hyperbolic tangent-like functions, which we refer to as tanh smoothing.

6.2.1. 1D smooth wave propagation

To assess numerical dissipation in smooth flow, we consider the propagation of a small-amplitude sinusoidal velocity perturbation with periodic boundary conditions on $[0, 1]$. The initial condition is $\rho = 1$, $u = \beta \sin(2\pi kx)$ with amplitude $\beta = 0.01$ and wavenumber $k = 20$, and $p = 1.6$ with $\gamma = 1.4$. The final time is $T = 0.4$.

We compare IGR ($\alpha = 2\Delta^2$), LAD, characteristic WENO, and component WENO at resolutions $m = 150$ and $m = 300$, using a characteristic WENO solution at $m = 10^4$ as the reference. Results are shown in Fig. 5: panels (a) and (b) compare all methods at $m = 150$ and $m = 300$, respectively; panel (c) shows the effect of increasing α on both IGR and LAD at $m = 300$, with $\alpha = 10\Delta^2$ and $100\Delta^2$. At the default $\alpha = 2\Delta^2$ and $m = 150$, IGR closely tracks the reference solution and is significantly less dissipative than either characteristic WENO or LAD. At $m = 300$, IGR achieves accuracy comparable to characteristic WENO, while LAD remains visibly more dissipative. In panel (c), increasing α from $10\Delta^2$ to $100\Delta^2$ has a negligible effect on IGR, whereas LAD’s amplitude at $\alpha = 100\Delta^2$ drops to roughly one-third of its default value. This asymmetry shows that IGR is notably more robust to the choice of α than LAD.

6.2.2. 1D Shu–Osher interaction

The Shu–Osher problem [40] models the interaction of a Mach 3 shock with a sinusoidal entropy perturbation (at constant pressure). We solve a rescaled version on the domain $[0, 1]$ with $\gamma = 1.4$ and final time $t = 0.18$. The post-shock state is $(\rho, u, p) = (27/7, 4\sqrt{35}/9, 31/3)$ and the pre-shock region has density $\rho = 1 + 0.2 \sin(16\pi x)$ with $u = 0$ and $p = 1$. We compare IGR ($\alpha = 2\Delta^2$) with LAD, characteristic WENO, and component-wise WENO at resolutions $m = 200, 300$, and 400 . The reference solution is computed with characteristic WENO at $m = 5000$. The initial condition is tanh smoothed for IGR and LAD; WENO uses the sharp initial condition directly. Results are shown in Fig. 6: the top row shows the full density profile, and the bottom row magnifies the post-shock oscillatory region. In the oscillatory region, IGR and LAD preserve the high-frequency density fluctuations more accurately than the WENO methods.




Fig. 5. Sine wave propagation comparing velocity profiles at final time $T = 0.4$. (a) and (b) compare all methods at $m = 150$ and $m = 300$, respectively. (c) shows IGR and LAD at $m = 300$ with $\alpha = 10\Delta x^2$ and $100\Delta x^2$. The reference solution (dashed blue) is computed with characteristic WENO and $m = 10^4$.




Fig. 6. Shu–Osher problem comparing density profiles at resolutions $m = 200, 300$, and 400 . The top row (a–c) shows the full domain; the bottom row (d–f) shows the post-shock oscillatory region. All methods are compared against a reference solution computed with characteristic WENO at $m = 5000$.

6.2.3. 1D Sod shock tube

The Sod shock tube [41] is a standard Riemann problem that produces a right-traveling shock, a contact discontinuity, and a left-traveling rarefaction. We solve it on the domain $[0, 1]$ with $\gamma = 1.4$ and final time $t = 0.2$. The initial condition is $(\rho, u, p) = (1, 0, 1)$ for $x < 0.5$ and $(\rho, u, p) = (0.125, 0, 0.1)$ for $x \geq 0.5$.

We compare IGR ($\alpha = 2\Delta x^2$) with LAD, characteristic WENO, and component WENO at resolutions $m = 200, 400$, and 800 . The initial condition is tanh smoothed for IGR and LAD with smoothing parameter $\epsilon = 4/m$, so that ϵ halves as m doubles, mirroring the $\sqrt{\alpha} \propto \Delta x$ scaling from Fig. 4. WENO uses the sharp initial condition directly. All results are compared against the exact Riemann solver solution.

Fig. 7 is organized as a 4×3 grid: the first two rows show density profiles (full and zoomed near the contact discontinuity), and the last two rows show specific internal energy profiles (full and zoomed). The zoomed panels focus on the contact discontinuity region ($x \in [0.65, 0.73]$), where differences between methods are most apparent. All methods converge at approximately first order, consistent with the presence of discontinuities. Near the contact, characteristic WENO achieves a sharper profile than IGR or LAD, owing to its use of the sharp initial condition. IGR and LAD produce profiles of similar width, set by the smoothing parameter $\epsilon \propto \Delta x$ of the initial condition rather than by numerical diffusion from the LF solver; the contact sharpens at the same rate as ϵ as resolution increases.




Fig. 7. Sod shock tube convergence study. Rows 1–2 (a–f): density profiles, full and zoomed near the contact discontinuity. Rows 3–4 (g–l): specific internal energy profiles, full and zoomed. Columns correspond to $(m, \varepsilon) = (200, 0.02)$, $(400, 0.01)$, and $(800, 0.005)$. WENO uses the sharp initial condition; LAD and IGR use smoothed initial conditions. All results are compared against the exact Riemann solver solution.

6.2.4. 1D Leblanc shock tube

The Leblanc shock tube [42] features a large density ratio (10^3) and pressure ratio ($\sim 10^8$), making it a stringent test of robustness under large jumps of the state variables.

We solve this problem on the domain $[0, 9]$ with $\gamma = 5/3$ and final time $t = 6$. The left (high-energy) state $(\rho, u, p) = (1, 0, 1/15)$ is initialized in a single cell at the left boundary ($x_0 = 0$), and the ambient right state $(\rho, u, p) = (10^{-3}, 0, (2/3) \times 10^{-9})$ fills the remainder of the domain.

We compare IGR ($\alpha = 2\Delta^2$), LAD, and characteristic WENO with $m = 450, 900$, and 1800 grid points across the 1D domain. All results are compared against the exact Riemann solver solution computed from the sharp initial condition. IGR and LAD require smoothed initial conditions and are thus only shown for $\varepsilon > 0$; WENO uses the sharp initial condition and appears in all panels. Here, ε is a dimensionless parameter controlling the width over which tanh smoothing is applied to the initial discontinuity.

Fig. 8 presents a convergence study in specific internal energy, organized as a grid with rows corresponding to different smoothing parameters ($\varepsilon = 0, 0.025, 0.05, 0.1$, amounting to different degrees of regularization of the initial condition) and columns showing different resolutions ($m = 450, 900, 1800$). As ε decreases and resolution increases, the IGR and LAD solutions converge toward the exact solution, demonstrating that the smoothed-IC approach recovers the correct entropy solution in the limit. However, IGR and LAD require progressively finer grid resolutions to avoid encountering negative densities or energies as the smoothing is decreased. Component-wise WENO crashes at each fixed ε and resolution, whereas Characteristic WENO achieves lower error than both IGR and LAD. The relative performance of IGR and LAD is close across all configurations, with IGR showing marginally lower errors in most cases.

6.2.5. 2D Riemann interaction with entropy-wave content

We consider a 2D Riemann problem [43] on a computational domain with sizes $[0, 1] \times [0, 1]$ and a grid resolution $N = 500^2$. The density, velocities, and pressure (ρ, u, v, p) in all regions of the computational domain are given in Eq. (30).

$$(\rho, u, v, p) = \begin{cases} (0.138, 1.206, 1.206, 0.029) & x \leq 0.75, y \leq 0.75 \\ (0.532, 0.000, 1.206, 0.300) & x \geq 0.75, y \leq 0.75 \\ (0.532, 1.206, 0.000, 0.300) & x \leq 0.75, y \geq 0.75 \\ (1.500, 0.000, 0.000, 1.500) & x \geq 0.75, y \geq 0.75 \end{cases} \quad (30)$$

An entropy wave is then introduced to this problem by adding fine-scale sinusoidal fluctuations in density (and consequently energy), while maintaining constant pressure. The resulting perturbed Riemann problem is run with IGR at $\alpha = 2\Delta^2$, where Δ is the grid spacing. For baseline comparisons, we use a 5th-order accurate WENO scheme with both the HLLC and Lax–Friedrichs (LF) Riemann solvers. We also run a well-resolved simulation at $N = 5000^2$ using WENO reconstruction and the LF Riemann solver, which we treat as the true solution in the vanishing viscosity limit. We apply tanh smoothing to the Riemann problem in Eq. (30), which is necessary for a stable IGR solution. Fig. 9 shows density heat maps for all numerical methods considered. Compared to the reference solution in Fig. 9(d), the flow structures match most closely with WENO5/HLLC (Fig. 9(c)), followed by IGR (Fig. 9(a)), and WENO5/LF (Fig. 9(b)). With the Lax–Friedrichs Riemann solver, WENO reconstruction completely dissipates the density fluctuations, while IGR preserves them. Also, WENO reconstruction introduces spurious grid-aligned artifacts near the shock boundary with both HLLC and LF Riemann solvers, which are absent with IGR.

6.2.6. 2D Double Mach reflection

The double Mach reflection problem [44] is a canonical test case for compressible flows. A planar shock passing through a reflecting wedge produces a shock reflection, yielding a self-similar flow parametrized by the Mach number and the slope angle.

The double-Mach case is set up in the usual way. We use a rectangular domain of size $[0, 4] \times [0, 1]$, with a Mach 10 oblique shock angled at 60° and a reflecting wall at the bottom of the domain. The reflective wall extends from $x = 1/6$ to $x = 4$, with the region to the left of $x = 1/6$ assigned post-shock flow. The top boundary follows the motion of the Mach 10 shock, with the shock starting at $x = 1/6$ at the bottom boundary. The left boundary has initial post-shock values, and the right wall has zero-gradient boundary conditions.

We conduct our simulation on an 800×200 grid, using IGR with $\alpha = 2\Delta^2$, where Δ is the grid size. Numerical stability requires tanh smoothing of the initial condition. Due to the reflective boundary condition, a low-order accurate reconstruction of the variables is necessary to maintain stability; for this purpose, we use third-order interpolation. Although WENO and IGR schemes are unstable with fifth-order accurate reconstruction, IGR is still stable with an increase in α to $10\Delta^2$.

The density contours at four simulation times with IGR are shown in Figs. 10(a)–10(d). The resulting flow structures match the expected results, except for an artifact originating at the shock's initial position, which becomes less prevalent as the simulation progresses and is largely absent at $T = 0.2$ in Fig. 10(d). This artifact arises because the initial condition only approximately matches the profile of a fully developed IGR shock. As it forms an IGR shock at the onset of the simulation, a wave is reflected, causing the artifact.




Fig. 8. Leblanc shock tube convergence study showing specific internal energy profiles. Each row corresponds to a different smoothing parameter ($\varepsilon = 0, 0.025, 0.05, 0.1$) and each column shows a different resolution ($m = 450, 900, 1800$). Runs not shown amount to simulations that break down due to negative densities or energies. All results are compared against the exact Riemann solver solution for the sharp initial condition. For $\varepsilon = 0$, only WENO is shown, as IGR and LAD require smoothed initial conditions. Component-wise WENO broke down in all of the above settings.

6.2.7. 2D Isentropic vortex

The isentropic vortex problem [45] is often used as a test problem because its exact solution enables evaluation of the accuracy and diffusion of numerical methods.

The vortex is initialized at the center of a periodic box of length $L = 10$. The initial conditions for density (ρ), velocity (u, v), and pressure (p) are given by

$$\begin{aligned}
 u &= u_{\infty} + y(S/2\pi) \exp(0.5\beta(1-r^2)) \\
 v &= v_{\infty} - x(S/2\pi) \exp(0.5\beta(1-r^2)) \\
 \rho &= \rho_{\infty} \left[1 - \frac{\rho_{\infty}}{p_{\infty}} \frac{S^2(\gamma-1)}{8\gamma\pi^2} \exp(\beta(1-r^2)) \right]^{\frac{1}{\gamma-1}} \\
 p &= p_{\infty} \left[1 - \frac{\rho_{\infty}}{p_{\infty}} \frac{S^2(\gamma-1)}{8\gamma\pi^2} \exp(\beta(1-r^2)) \right]^{\frac{\gamma}{\gamma-1}}
 \end{aligned} \tag{31}$$

We use vortex strength $S = 5$, $\beta = 1$, and $\gamma = 1.4$ in Eq. (31), and $r^2 = x^2 + y^2$. The free stream values for density (ρ_{∞}) and pressure (p_{∞}) are unity, and the x and y velocities are $u_{\infty} = 0.1$ and $v_{\infty} = 0$. The analytical solution for this problem advects the vortex at




Fig. 9. Riemann test problem with $N = 500^2$ grid points using (a) IGR/LF, (b) WENO5/LF, and (c) WENO5/HLLC with a (d) reference vanishing viscosity solution using WENO5/LF at $N = 5000^2$ grid points.




Fig. 10. Double Mach problem with $m_x \times m_y = 800 \times 200$ grid points at (a) $T = 0.05$, (b) $T = 0.10$, (c) $T = 0.15$, and (d) $T = 0.20$ using IGR and tanh smoothing.

velocity $u_\infty = 0.1$ in the x direction of the periodic domain. We evaluate the solution error after a fixed number of such periodic pass-through periods and compare the solution to the initial condition.

We conduct an error convergence test using 5th-order accurate IGR and WENO schemes. IGR tests are conducted at varying values of $\alpha = (0, 2, 10)\Delta^2$, with Δ being the grid spacing. The error $\|e\|_\infty$ in pressure p after 4 periods at varying grid sizes ($N = 50^2$, 100^2 , 150^2 , and 200^2) is evaluated and plotted in Fig. 11. Both the WENO5 and IGR with $\alpha = 0$ exhibit 4th-order accurate error convergence in pressure, as expected. Due to the discretization of the right-hand side of the elliptic solve for Σ (Eq. (12)), IGR with $\alpha > 0$ can be at most 2nd-order accurate. The error grows with increasing α due to increased smoothing at a fixed resolution.

To examine the effects of numerical dissipation, we investigate the error growth for the aforementioned schemes as the number of periods across the computational domain increases. We conduct WENO simulations with Lax–Friedrichs (LF) and HLLC Riemann




Fig. 11. Convergence rate for error $\|e\|_\infty$ in pressure with grid size for IGR with varying α and WENO5.




Fig. 12. Growth of error $\|e\|_\infty$ in pressure for an advecting isentropic vortex with grid sizes (a) $N = 200^2$ and (b) $N = 400^2$ versus the number of periods across the domain.

solvers and IGR with LF. The test is conducted over 30 periods at two different grid resolutions ((a) $N = 200^2$ and (b) $N = 400^2$) with the growth of error $\|e\|_\infty$ in pressure p shown in Fig. 12.

One sees that the error associated with projecting the initial condition stabilizes after sufficient T . After the error becomes consistent, its growth rate exhibits two primary, well-established [46] regimes: an initial linear growth phase and a later exponential one. The linear regime, which arises from numerical dissipation, is prominent for WENO (both Riemann solvers) and IGR with $\alpha = 0$. In Fig. 12(a), this lasts up to $T = 5$ and 20 periods for WENO with LF and HLLC Riemann solvers, respectively, and up to $T = 25$ periods for IGR with $\alpha = 0$. This regime is absent for IGR with $\alpha > 0$, with the error levels remaining constant due to lower numerical dissipation.

All simulations eventually exhibit an exponential regime due to a destabilized vortex, leading to exponential error growth. This instability is independent of the numerical method and arises from the problem definition and its accompanying boundary conditions [46]. The onset of the exponential regime, however, occurs earlier for the WENO schemes: $T = 5$ and 20 periods with LF and HLLC in Fig. 12(a). We observe this onset for IGR at $T = 21$ and 25 periods with $\alpha = 10\Delta^2$ and $2\Delta^2$ in Fig. 12(a), with a smaller value of α further postponing it.

6.2.8. 3D Taylor–Green vortex

The 3D Taylor–Green vortex (TGV) [47] tests IGR’s ability to represent compressible turbulence. The TGV is set up in a 3D cube with side lengths $L = 2\pi$, with air at uniform density ($\rho = 1$) and the initial conditions for velocities (u, v, w) and pressure p given by Eq. (32), where p_0 is the reference (ambient) pressure.

$$\begin{aligned}
 u &= V_0 \sin(x) \cos(y) \sin(z) \\
 v &= -V_0 \cos(x) \sin(y) \sin(z) \\
 w &= 0 \\
 p &= p_0 + \frac{V_0^2}{16} (\cos(2x) + \cos(2y))(\cos(2z) + 2)
 \end{aligned} \tag{32}$$

The Mach number V_0/c_0 is 1.25, where $c_0 = \sqrt{\gamma p_0/\rho}$ is the reference speed of sound, so shocklets appear in the flow [48]. The Reynolds number is $\text{Re} = 10^4$ to have a sufficiently long inertial regime in the energy cascade, where the turbulent kinetic energy $E(k) \propto k^{-5/3}$. The inertial regime for this case typically ends at wave number $k = 2\pi/(60\eta)$ [49], where $\eta \propto \text{Re}^{-3/4}$ is the Kolmogorov scaling. For $\text{Re} = 10^4$, this corresponds to an inertial regime that lasts up to $k = 100$. We thus conduct simulations at $N = m_x m_y m_z = 128^3, 256^3$, and 512^3 , which are sufficiently high to represent the inertial regime.




Fig. 13. Turbulent kinetic energy cascade for IGR/LF, WENO5/HLLC, and WENO5/LF for $N = 128^3$, 256^3 , and 512^3 grid points, normalized by a WENO5/HLLC solution that used 2048^3 grid points.

We use IGR with $\alpha = 2\Delta^2$ and the Lax–Friedrichs (LF) Riemann solver, and compare the results to a baseline 5th-order accurate WENO (WENO5) scheme with HLLC and Lax–Friedrichs Riemann solvers. The turbulent energy cascade is then calculated using fast Fourier transforms and plotted at $t_c = 20L/V_0$. The results are time-averaged across $t = 0.2t_c$ to minimize spurious oscillations. The energy cascade is normalized with a ground truth solution using WENO5 and HLLC at $N = 2048^3$ and plotted in Fig. 13. The results indicate notably greater dissipation with WENO than with IGR for the same Riemann solver (LF) and grid resolution, where IGR can sustain the inertial regime over a broader range of wavenumbers. WENO reconstruction with HLLC, however, can improve upon IGR, with the gap narrowing at higher resolutions ($N = 512^3$).

6.2.9. 3D multi-jet configuration

Many-engine rockets, such as the SpaceX Super Heavy, are powered by many smaller engines than the industry standard at the time of writing. This choice comes with a key engineering challenge: high-speed exhaust plumes interact, leading to upstream-going waves and base heating [50].

We simulate an array of 33 rocket engines with an outlet velocity of Mach 10 in atmospheric conditions, a configuration inspired by the SpaceX Super Heavy. The simulation uses 600 grid cells per engine outlet, resulting in 3.3×10^{12} grid cells on a rectilinear grid. This configuration was also used by the authors in the Gordon Bell work of Wilfong et al. [24].

We use IGR with $\alpha = 2\Delta^2$, where Δ is the grid spacing. We apply Gaussian smoothing to the initial condition to ensure numerical stability. The velocity contours for this simulation are shown in Fig. 14, and the fine-scale structures of interacting exhaust plumes are captured using IGR. A comparison of exhaust plumes for IGR and the baseline WENO simulation with HLLC Riemann solver for a simpler three-jet configuration is also presented in Fig. 14. Due to the baseline simulation's grid-dependent nature, grid-aligned artifacts are more prominent than with IGR.

7. Comparison to related work, discussion, and conclusions

7.1. Comparison to related work

Upwind finite-volume methods trace back to Godunov [51], who solved local Riemann problems at cell interfaces to obtain numerical fluxes. The piecewise parabolic method (PPM) [52] and WENO [1,16] extended this to higher spatial order. More complete Riemann solvers were developed to reduce interface dissipation, including the Roe [53], HLLC [35,36], and AUSM+ [54] schemes. High-order linear reconstruction near shocks introduces Gibbs oscillations; limiters such as MUSCL [30,31] and ENO [29] suppress these by adaptively reducing the approximation order near discontinuities. The resulting mismatch between left- and right-interface states then determines the amount of dissipation added by the Riemann solver. Flux-corrected transport (FCT) [55,56], multi-dimensional optimal order detection (MOOD) [57], and convex limiting [18,19] take an a posteriori approach, applying corrections after the fact to maintain physical bounds. Entropy-stable schemes [58] and central schemes [59] provide dissipation control without requiring a full Riemann solve. The Runge–Kutta discontinuous Galerkin (DG) framework [60] achieves high-order accuracy on unstructured meshes but still requires limiting or artificial viscosity near discontinuities.

Each of these approaches introduces nonlinear cell-to-cell coupling, thereby dissipating fine-scale features. Approximate Riemann solvers such as Roe's scheme can additionally produce carbuncle instabilities on grid-aligned shocks [37,61]. IGR produces smooth PDE solutions, so a linear Lax–Friedrichs flux suffices without a limiter or Riemann solver, avoiding these issues.

The 2D Riemann results show that IGR better preserves small-scale density fluctuations than WENO/LF, and avoids the grid-aligned artifacts present in both WENO/LF and WENO/HLLC. The isentropic vortex results show that IGR has a lower error growth




Fig. 14. Application of IGR to Mach 10 rocket simulations with strong shocks. Panel (a) shows an artistic rendering of a 33-jet configuration. Panels (b) and (c) show a comparison between IGR and WENO simulations of a simpler three-jet configuration. The IGR result in (b) avoids the grid-aligned artifacts present in the (c) WENO case. Images adapted from Wilfong et al. [24] with author permission.

rate than WENO with either solver, reflecting reduced numerical dissipation in smooth flow. The Taylor–Green vortex results show that IGR sustains the turbulent inertial range longer than WENO/LF at the same resolution.

Viscous regularizations, including artificial viscosity [4], entropy viscosity [10], LAD [6,8], and PDE-based AV [9], smooth shocks through added dissipation, which also damps turbulence and acoustic waves. IGR is an inviscid regularization; the Shu–Osher results show better preservation of high-frequency wave content than LAD, and the Taylor–Green vortex results show lower dissipation than WENO/LF at the same resolution. Shock tracking [62] avoids smearing by explicitly evolving the shock surface, recovering sharp profiles without diffusion. This approach, however, is more complex to implement for multi-dimensional and topology-changing configurations. IGR produces smooth shock profiles of controllable width $\sqrt{\alpha}$ within a standard finite-volume solver.

7.2. Limitations of this work

IGR requires initial conditions that are smooth on the scale $\sqrt{\alpha}$. Sharp initial data must be preprocessed with a smoothing step, as done in several test cases here, but this has no physical consequence since IGR solutions converge to the correct entropy solution as $\alpha \rightarrow 0$. This is a practical consideration when setting up a simulation.

The comparisons in this work also show that WENO5/HLLC can match or outperform IGR in certain regimes. In the 2D Riemann test, WENO5/HLLC reproduces the reference solution most closely overall, and in the Taylor–Green vortex, the performance gap between IGR and WENO5/HLLC narrows at higher resolutions. IGR’s advantages are most pronounced when paired with a more dissipative (and cheaper) Riemann solver such as Lax–Friedrichs, or when fine-scale features must be preserved alongside shocks. In this work, we use combine IGR with the linear Lax–Friedrichs Riemann solver and a linear fifth-order reconstruction obtained by WENO5 in smooth regions. However, we emphasize that the use of IGR is mostly agnostic to the choice of numerical method. We believe that the combination of IGR with more accurate high-order discretizations, such as minimally dispersive methods developed for turbulence simulation [63], holds great promise that we aim to explore in future work.

The convergence order of a numerical method is not necessarily reflective of its performance in the presence of localized features. More important is the cost–resolution tradeoff: “how much computation is needed to resolve flow features of a given scale”. For this goal, increasing the order of a numerical method is helpful only to a point [64]. The present work focuses on WENO5 due to its widespread use in the CFD community. We point out that alternatives such as piecewise linear (PLM) and piecewise quadratic (PPM) reconstructions were reported to yield better cost–accuracy tradeoffs by [65]. This is subject to the caveat that the ever-widening gap between memory-access and compute costs pushes the balance toward higher-order methods that are more arithmetically intensive. Nevertheless, we emphasize that IGR is agnostic to the underlying numerical method and that this work serves as a template for deploying IGR to many other numerical methods. Just like in the case of WENO, IGR could allow deploying PLM and PPM without slope limiters. This decreases computational cost as well as the diffusive effect of the approximate Riemann solver. We defer an investigation of this setting to future work. The under-resolved cases presented here, such as the rocket plume, operate in an implicit LES regime in which the numerical scheme and IGR together act as the effective filter. The application of IGR in explicit LES using, for instance, a Smagorinsky-type subgrid model, is an interesting direction of future work.

We use $\alpha = 2\Delta^2$ throughout, but we set it empirically. Rigorous guidance on the optimal choice of α for a given problem does not yet exist.

Because the entropic pressure Σ is defined by a linear, self-adjoint elliptic operator, differentiating through the IGR system is structurally well-posed: the adjoint requires only a second elliptic solve of the same type. Adjoint-based sensitivity computation with IGR has not been demonstrated in this work and is identified as a promising direction for future research. The theoretical analysis

of IGR convergence to entropy solutions has only been established for the 1D pressureless case [25]; extending these guarantees to the full compressible Navier–Stokes system remains open.

All test cases in this work use an ideal gas equation of state. The behavior of IGR in multi-species, real-gas, or reactive thermodynamic systems has not been investigated.

7.3. Conclusion

This work demonstrates the utility of information geometric regularization (IGR) within a finite-volume framework for the compressible Navier–Stokes equations. IGR replaces singular shocks with smooth profiles of user-controlled width $\sqrt{\alpha}$, eliminating the need for limiters and nonlinear Riemann solvers. The resulting solver reduces to a linear reconstruction paired with a Lax–Friedrichs solve, lowering computational overhead and enabling kernel fusion relative to standard limiter-based approaches.

The 1–3D test suite spans shock tubes, the Shu–Osher interaction, a 2D Riemann problem with entropy-wave content, the double Mach reflection, and the Taylor–Green vortex. Across all cases, IGR preserves fine-scale flow features that limiter-based and viscous methods dissipate, and does so with a simpler numerical scheme. These properties suggest that IGR is a viable replacement for conventional shock-capturing in finite-volume codes, particularly in compressible flows where shocks coexist with fine-scale wave content.

CRediT authorship contribution statement

Anand Radhakrishnan: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Benjamin Wilfong:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Spencer H. Bryngelson:** Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Florian Schäfer:** Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process.

During the preparation of this work, the author(s) used Claude Code to set up numerical experiments and tabulate results for figure generation. Grammarly was used to refine some of the prose. After using these tools/services, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content of the published article.

Declaration of competing interest

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

Acknowledgments

This work was partially supported by the Predictive Science Academic Alliance Program (PSAAP Award DE-NA0004261 - “The Center for Information Geometric Mechanics and Optimization (CIGMO)”) managed by the NNSA (National Nuclear Security Administration) Office of Advanced Simulation.

FS gratefully acknowledges support from the Air Force Office of Scientific Research under award number FA9550-23-1-0668 (Information Geometric Regularization for Simulation and Optimization of Supersonic Flow) and from the Alfred P. Sloan Foundation via a Sloan Research Fellowship in Mathematics.

SHB acknowledges support for this work from NVIDIA Corporation and Advanced Micro Devices, Inc. (AMD) through hardware donations.

This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725, under allocation CFD154.

This work used Delta at the National Center for Supercomputing Applications through allocation PHY210084 from the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) program, which is supported by National Science Foundation grants #2138259, #2138286, #2138307, #2137603, and #2138296. This research is part of the Delta research computing project, which is supported by the National Science Foundation (award OCI 2005572), and the State of Illinois. Delta is a joint effort of the University of Illinois at Urbana–Champaign and its National Center for Supercomputing Applications. We thank the anonymous reviewers for their valuable comments and suggestions that helped us to improve this manuscript.

Data availability

All code is available at <https://github.com/sbryngelson/figr> (multidimensional results using MFC) and https://github.com/f-t-s/IGR_Manual_1dFVM (unidimensional results in Julia).

References

- [1] C.-W. Shu, Essentially non-oscillatory and weighted essentially non-oscillatory schemes for hyperbolic conservation laws, in: B. Cockburn, C. Johnson, C.-W. Shu, E. Tadmor (Eds.), *Advanced Numerical Approximation of Nonlinear Hyperbolic Equations*, in: *Lecture Notes in Mathematics*, vol. 1697, Springer, 1998, pp. 325–432.
- [2] D. Gottlieb, C.-W. Shu, On the Gibbs phenomenon and its resolution, *SIAM Rev.* 39 (4) (1997) 644–668.
- [3] C. Canuto, M. Hussaini, A. Quarteroni, T.A. Zang, *Spectral Methods: Fundamentals in Single Domains*, Springer, 2006.
- [4] J. VonNeumann, R.D. Richtmyer, A method for the numerical calculation of hydrodynamic shocks, *J. Appl. Phys.* 21 (3) (1950) 232–237.
- [5] G. Puppo, Numerical entropy production for central schemes, *SIAM J. Sci. Comput.* 25 (4) (2004) 1382–1415.
- [6] A.W. Cook, W.H. Cabot, Hyperviscosity for shock-turbulence interactions, *J. Comput. Phys.* 203 (2) (2005) 379–385.
- [7] B. Fiorina, S.K. Lele, An artificial nonlinear diffusivity method for supersonic reacting flows with shocks, *J. Comput. Phys.* 222 (1) (2007) 246–264.
- [8] A. Mani, J. Larsson, P. Moin, Suitability of artificial bulk viscosity for large-eddy simulation of turbulent flows with shocks, *J. Comput. Phys.* 228 (19) (2009) 7368–7374.
- [9] G.E. Barter, D.L. Darmofal, Shock capturing with PDE-based artificial viscosity for DGFEM: Part I. Formulation, *J. Comput. Phys.* 229 (5) (2010) 1810–1827.
- [10] J.-L. Guermond, R. Pasquetti, B. Popov, Entropy viscosity method for nonlinear conservation laws, *J. Comput. Phys.* 230 (11) (2011) 4248–4267.
- [11] O.P. Bruno, J.S. Hesthaven, D.V. Leibovici, FC-based shock-dynamics solver with neural-network localized artificial-viscosity assignment, *J. Comput. Phys.* X 15 (2022) 100110.
- [12] V. Dolejší, M. Feistauer, C. Schwab, On some aspects of the discontinuous Galerkin finite element method for conservation laws, *Math. Comput. Simulation* 61 (3–6) (2003) 333–346.
- [13] J. Glaubitz, P. Öffner, H. Ranocha, T. Sonar, Smooth and compactly supported viscous sub-cell shock capturing for discontinuous Galerkin methods, *J. Sci. Comput.* 79 (1) (2019) 249–272.
- [14] B. Van Leer, Towards the ultimate conservative difference scheme. V. A second-order sequel to Godunov's method, *J. Comput. Phys.* 32 (1) (1979) 101–136.
- [15] D. Ray, J.S. Hesthaven, An artificial neural network as a troubled-cell indicator, *J. Comput. Phys.* 367 (2018) 166–191.
- [16] X.-D. Liu, S. Osher, T. Chan, Weighted essentially non-oscillatory schemes, *J. Comput. Phys.* 115 (1) (1994) 200–212.
- [17] A. Harten, B. Engquist, S. Osher, S.R. Chakravarthy, Uniformly high order accurate essentially non-oscillatory schemes, III, *J. Comput. Phys.* 131 (1) (1997) 3–47.
- [18] J.-L. Guermond, M. Nazzari, B. Popov, I. Tomas, Second-order invariant domain preserving approximation of the Euler equations using convex limiting, *SIAM J. Sci. Comput.* 40 (5) (2018) A3211–A3239.
- [19] D. Kuzmin, Monolithic convex limiting for continuous finite element discretizations of hyperbolic conservation laws, *Comput. Methods Appl. Mech. Engrg.* 361 (2020) 112804.
- [20] C. Lozano, Watch your adjoints! lack of mesh convergence in inviscid adjoint solutions, *AIAA J.* 57 (9) (2019) 3991–4006.
- [21] D. Bodony, A. Fikl, Adjoint-based sensitivity of shock-laden flows, in: *Proceedings of the 2022 CTR summer program*, 2022.
- [22] M. Nemec, M. Aftosis, Aerodynamic shape optimization using a cartesian adjoint method and CAD geometry, in: *24th AIAA Applied Aerodynamics Conference*, 2006, p. 3456.
- [23] R. Cao, F. Schäfer, Information geometric regularization of the barotropic Euler equation, 2023, arXiv preprint arXiv:2308.14127.
- [24] B. Wilfong, A. Radhakrishnan, H. Le Berre, D.J. Vickers, T. Prathi, N. Tselepidis, B. Dorschner, R. Budiardja, B. Cornille, S. Abbott, F. Schäfer, S.H. Bryngelson, Simulating many-engine spacecraft: Exceeding 1 quadrillion degrees of freedom via information geometric regularization, in: *Proceedings of SC '25: The International Conference for High Performance Computing, Networking, Storage and Analysis*, 2025, pp. 14–24.
- [25] R. Cao, F. Schäfer, Information geometric regularization of unidimensional pressureless Euler equations yields global strong solutions, 2024, arXiv preprint arXiv:2411.15121.
- [26] W. Barham, B.K. Tran, B.S. Southworth, F. Schäfer, Hamiltonian information geometric regularization of the compressible Euler equations, 2025, arXiv preprint arXiv:2512.13948.
- [27] B. Guelmame, D. Clamond, S. Junca, Hamiltonian regularisation of the unidimensional barotropic Euler equations, *Nonlinear Anal. Real World Appl.* 64 (2022) 103455.
- [28] J. Koreeda, H. Yanagisawa, K. Maeno, H. Honma, S. Bystrov, V. Ivanov, F. Shugayev, Front structures of strong shock waves in air, in: *Shock Waves @ Marseille II: Physico-Chemical Processes and Nonequilibrium Flow*, Springer, 1995, pp. 263–268.
- [29] A. Harten, B. Engquist, S. Osher, S.R. Chakravarthy, Uniformly high order accurate essentially non-oscillatory schemes, III, *J. Comput. Phys.* 71 (2) (1987) 231–303.
- [30] B. van Leer, Towards the ultimate conservative difference scheme. V. A second-order sequel to Godunov's method, *J. Comput. Phys.* 32 (1) (1979) 101–136.
- [31] B. van Leer, Towards the ultimate conservative difference scheme. I. The quest of monotonicity, *Lecture Notes in Phys.* 18 (1973) 163–168.
- [32] V. Arnold, Sur la géométrie différentielle des groupes de Lie de dimension infinie et ses applications à l'hydrodynamique des fluides parfaits, *Ann. L'Institut Fourier* 16 (1) (1966) 319–361.
- [33] B. Khesin, G. Misiolek, K. Modin, Geometric hydrodynamics and infinite-dimensional Newton's equations, *Bull. Am. Math. Soc.* 58 (3) (2021) 377–442.
- [34] A. Harten, P.D. Lax, B. van Leer, On upstream differencing and Godunov-type schemes for hyperbolic conservation laws, *SIAM Rev.* 25 (1) (1983) 35–61.
- [35] E.F. Toro, M. Spruce, W. Speares, Restoration of the contact surface in the HLL-Riemann solver, *Shock Waves* 4 (1) (1994) 25–34.
- [36] E.F. Toro, *Riemann Solvers and Numerical Methods for Fluid Dynamics: A Practical Introduction*, third ed., Springer, 2009.
- [37] Y. Chauvat, J.-M. Moschetta, J. Gressier, Shock wave numerical structure and the carbuncle phenomenon, *Internat. J. Numer. Methods Fluids* 47 (8–9) (2005) 903–909.
- [38] S. Gottlieb, C.-W. Shu, E. Tadmor, Strong stability-preserving high-order time discretization methods, *SIAM Rev.* 43 (1) (2001) 89–112.
- [39] P.S. Rawat, X. Zhong, On high-order shock-fitting and front-tracking schemes for numerical simulation of shock-disturbance interactions, *J. Comput. Phys.* 229 (19) (2010) 6744–6780.
- [40] C.-W. Shu, S. Osher, Efficient implementation of essentially non-oscillatory shock-capturing schemes, *J. Comput. Phys.* 77 (2) (1988) 439–471.
- [41] G.A. Sod, A survey of several finite difference methods for systems of nonlinear hyperbolic conservation laws, *J. Comput. Phys.* 27 (1) (1978) 1–31.
- [42] R. Loubère, M.J. Shashkov, A subcell remapping method on staggered polygonal grids for arbitrary Lagrangian–Eulerian methods, *J. Comput. Phys.* 209 (1) (2005) 105–138.
- [43] T. Liang, L. Fu, A new high-order shock-capturing TENO scheme combined with skew-symmetric-splitting method for compressible gas dynamics and turbulence simulation, *Comput. Phys. Comm.* 302 (2024) 109236.
- [44] P. Woodward, P. Colella, The numerical simulation of two-dimensional fluid flow with strong shocks, *J. Comput. Phys.* 54 (1) (1984) 115–173.
- [45] C.-W. Shu, Essentially non-oscillatory and weighted essentially non-oscillatory schemes for hyperbolic conservation laws, in: *Advanced Numerical Approximation of Nonlinear Hyperbolic Equations: Lectures Given At the 2nd Session of the Centro Internazionale Matematico Estivo (CIME) Held in Cetraro, Italy, June 23–28, 1997*, Springer, 2006, pp. 325–432.
- [46] S.C. Spiegel, H. Huynh, J.R. DeBonis, A survey of the isentropic Euler vortex problem using high-order methods, in: *22nd AIAA Computational Fluid Dynamics Conference*, 2015, p. 2444.

- [47] M.E. Brachet, D.I. Meiron, S.A. Orszag, B.G. Nickel, R.H. Morf, U. Frisch, Small-scale structure of the Taylor–Green vortex, *J. Fluid Mech.* 130 (1983) 411–452.
- [48] D.J. Lusher, N.D. Sandham, Assessment of low-dissipative shock-capturing schemes for the compressible Taylor–Green vortex, *AIAA J.* 59 (2) (2021) 533–545.
- [49] S.B. Pope, *Turbulent Flows*, Cambridge University Press, Cambridge, UK, 2000.
- [50] M. Mehta, F. Canabal, S.B. Tashakkor, S.D. Smith, Numerical base heating sensitivity study for a four-rocket engine core configuration, *J. Spacecr. Rockets* 50 (3) (2013) 509–526.
- [51] S.K. Godunov, A difference method for numerical calculation of discontinuous solutions of the equations of hydrodynamics, *Mat. Sb.* 47 (3) (1959) 271–306.
- [52] P. Colella, P.R. Woodward, The piecewise parabolic method (PPM) for gas-dynamical simulations, *J. Comput. Phys.* 54 (1) (1984) 174–201.
- [53] P.L. Roe, Approximate Riemann solvers, parameter vectors, and difference schemes, *J. Comput. Phys.* 43 (2) (1981) 357–372.
- [54] M.-S. Liou, A sequel to AUSM, Part II: AUSM+–up for all speeds, *J. Comput. Phys.* 214 (1) (2006) 137–170.
- [55] J.P. Boris, D.L. Book, Flux-corrected transport. I. SHASTA, a fluid transport algorithm that works, *J. Comput. Phys.* 11 (1) (1973) 38–69.
- [56] S.T. Zalesak, Fully multidimensional flux-corrected transport algorithms for fluids, *J. Comput. Phys.* 31 (3) (1979) 335–362.
- [57] S. Clain, S. Diot, R. Loubère, A high-order finite volume method for systems of conservation laws—Multi-dimensional optimal order detection (MOOD), *J. Comput. Phys.* 230 (10) (2011) 4028–4050.
- [58] E. Tadmor, The numerical viscosity of entropy stable schemes for systems of conservation laws. I, *Math. Comp.* 49 (179) (1987) 91–103.
- [59] A. Kurganov, E. Tadmor, New high-resolution central schemes for nonlinear conservation laws and convection–diffusion equations, *J. Comput. Phys.* 160 (1) (2000) 241–282.
- [60] B. Cockburn, C.-W. Shu, The Runge–Kutta discontinuous Galerkin method for conservation laws V: Multidimensional systems, *J. Comput. Phys.* 141 (2) (1998) 199–224.
- [61] J.J. Quirk, A contribution to the great Riemann solver debate, *Internat. J. Numer. Methods Fluids* 18 (6) (1994) 555–574.
- [62] J. Glimm, J.W. Grove, X. Li, K.-M. Shyue, Y. Zeng, Q. Zhang, Three-dimensional front tracking, *SIAM J. Sci. Comput.* 19 (3) (1998) 703–727.
- [63] S.K. Lele, Compact finite difference schemes with spectral-like resolution, *J. Comput. Phys.* 103 (1) (1992) 16–42.
- [64] W.J. Rider, J.R. Kamm, How effective are high-order approximations in shock-capturing methods? is there a law of diminishing returns? in: *Computational Fluid Dynamics 2004: Proceedings of the Third International Conference on Computational Fluid Dynamics, ICCFD3, Toronto, 12–16 July 2004*, Springer, 2006, pp. 401–405.
- [65] J. Greenough, W. Rider, A quantitative comparison of numerical methods for the compressible Euler equations: fifth-order WENO and piecewise-linear Godunov, *J. Comput. Phys.* 196 (1) (2004) 259–281.