

Assessment of low-dissipative shock-capturing schemes for the compressible Taylor-Green vortex

David J. Lusher^{*} and Neil D. Sandham[†]

University of Southampton, Aerodynamics and Flight Mechanics Group. Boldrewood Campus, Southampton, SO16 7QF, United Kingdom

Interactions between shock-waves and turbulence are ubiquitous in high-speed flows of practical aeronautical interest. Recent advances in computational power have made Implicit Large Eddy Simulation (ILES) and Direct Numerical Simulation (DNS) feasible tools for investigating the underlying physical mechanisms involved. However, numerical methods for shock-capturing introduce high levels of numerical dissipation to the whole flow-field, making them a poor choice for resolving the small scales of turbulence. In this work the efficacy of a selection of low-dissipative and hybrid Weighted and Targeted Essentially Non-Oscillatory (WENO/TENO) shock-capturing methods is assessed. An extension of the classic subsonic Taylor-Green vortex problem is presented up to Mach 1.25, where compressibility and dilatational dissipation become important. The presence of strong shock-waves is demonstrated, which merge and interact with one another to form complex shock systems. A supersonic test case is then specified, using a Reynolds number of 1600 to ensure a wide range of scales are present to test the numerical schemes. Low-dissipative TENO schemes are found to offer substantial improvements in resolution over established WENO methods for a comparable computational cost.

Nomenclature

$(x, y, z), (N_x, N_y, N_z)$	=	spatial coordinates, and number of grid points
$(\rho, u, v, w, p, T, \mu)$	=	non-dimensional density, velocities, pressure, temperature, and viscosity
Q	=	Q-Criterion
q_k, τ_{ik}	=	heat-flux and viscous stress tensor
$\Delta t, t$	=	time step, and simulation time
$\varepsilon^T, \varepsilon^S, \varepsilon^D$	=	total, solenoidal, and dilatational viscous dissipation
Re, Pr, M_{ref}	=	Reynolds number, Prandtl number, and reference Mach number

^{*}Postdoctoral Research Fellow, Aerodynamics and Flight Mechanics Group, University of Southampton.

[†]Professor of Aerospace Engineering, Aerodynamics and Flight Mechanics Group, University of Southampton.

I. Introduction

SIMULATION of the compressible Navier-Stokes equations at high Mach numbers requires numerical schemes that are capable of resolving discontinuities in the form of shock-waves. Within a finite-difference framework, standard central differencing generates spurious oscillations at discontinuities and becomes unstable, making it unsuitable for shocked flows [1]. To ensure stability a common approach is to add an artificial dissipation term to smear shocks over several grid points while trying to maintain the accuracy of the solution. Shock-capturing methods however introduce excessive levels of numerical dissipation to the flow-field, making them a poor choice for resolving small scale structures in transitional and fully turbulent flows. To obtain similar resolution to a non-dissipative scheme, excessively fine grids are required. These large grids are impractical for Implicit Large Eddy Simulation (ILES) or Direct Numerical Simulation (DNS), and severely limit the scale of the problem that can be tackled. Hybrid methods pair a non-dissipative scheme with shock-capturing applied only in areas of high gradients [2], reducing the numerical dissipation that is introduced into the flow field. While these methods have been widely used, care has to be taken with the choice of shock-sensor and tuning may be required between different physical problems.

In the past few decades considerable efforts have been made to reduce the numerical dissipation of the underlying shock-capturing schemes themselves, such as the improved WENO-Z scheme [3], [4] tested in the present work. Targeted Essentially non-Oscillatory (TENO) schemes [5], [6], aim to improve upon the existing WENO schemes via a discrete stencil selection method and a physically motivated scale separation weighting procedure. A further development [7] added a shock sensor to the TENO schemes to enable adaptive dissipation control in different regions of the computational domain. Previous studies investigating shock-capturing methods include [2], in which shock-fitting, WENO, hybrid-WENO, and artificial diffusivity (AD) schemes were compared for a suite of test problems. Standard WENO and AD schemes were found to be unsuitable for compressible turbulence, as the excessive numerical dissipation overwhelmed the physical dissipation and led to under-predictions of the upper 3/4 of resolvable wavenumbers. Hybrid-WENO methods fared better but there was a clear difficulty in creating a shock sensor applicable to all flow types. A more recent study [8] compared some improved WENO schemes to central-AD and localised artificial diffusivity (LAD). The WENO schemes proved more effective in robustly capturing shocks than the central-AD methods, albeit at a $\sim (2.5 - 3) \times$ increase in computational cost. LAD schemes performed better, but were poor compared to WENO for problems containing strong shocks. The aim of this work is to extend the current literature to include TENO schemes [5] and hybrid methods, for a supersonic problem containing both multiple shock-waves and compressible turbulence.

All of the schemes in this work are shock-capturing methods selected from the family of Essentially Non-Oscillatory (ENO) schemes. The schemes are assessed for a compressible, supersonic, version of the well-known Taylor-Green vortex (TGV) problem [9], which involves the laminar to turbulent transition of a decaying vortex. TGV has commonly been used at Mach 0.1 to assess numerical methods [10, 11] in both inviscid and viscous forms, despite the incompressible, smooth nature of the problem. The effect of an increasing Mach number is demonstrated in this work for the viscous TGV

problem up to Mach 1.25. Unlike the incompressible problem, multiple interacting shock-waves are shown to be present throughout the domain at Mach 1.25, making it a more appropriate test of shock-capturing schemes for compressible turbulence. After highlighting the main departures from the subsonic problem, the effectiveness of low-dissipative shock-capturing methods is assessed, in terms of both their flow resolving ability and relative computational cost. The compressible TGV problem at higher Mach numbers is a challenging test case for numerical schemes, but has not been widely explored in the literature. One notable exception is the work of [12], which investigated the evolution of vortex-surface fields in a compressible TGV. This study was performed at a low Reynolds number of $Re = 400$. A Reynolds number of $Re = 1600$ is selected in the present work to ensure that a wide range of scales are present to test the numerical schemes.

II. Numerical methods

A. Governing equations

The governing equations for all simulations in this work are the dimensionless compressible Navier-Stokes equations for a Newtonian fluid. Applying conservation of mass, momentum and energy in three spatial directions x_i ($i = 0, 1, 2$) results in a system of five partial differential equations given by

$$\frac{\partial \rho}{\partial t} + \frac{\partial}{\partial x_k} (\rho u_k) = 0, \quad (1)$$

$$\frac{\partial}{\partial t} (\rho u_i) + \frac{\partial}{\partial x_k} (\rho u_i u_k + p \delta_{ik} - \tau_{ik}) = 0, \quad (2)$$

$$\frac{\partial}{\partial t} (\rho E) + \frac{\partial}{\partial x_k} \left(\rho u_k \left(E + \frac{p}{\rho} \right) + q_k - u_i \tau_{ik} \right) = 0, \quad (3)$$

with Fourier's heat flux q_k and viscous stress tensor τ_{ij} defined as

$$q_k = \frac{-\mu}{(\gamma - 1) M_{\text{ref}}^2 Pr Re} \frac{\partial T}{\partial x_k}, \quad (4)$$

$$\tau_{ik} = \frac{\mu}{Re} \left(\frac{\partial u_i}{\partial x_k} + \frac{\partial u_k}{\partial x_i} - \frac{2}{3} \frac{\partial u_j}{\partial x_j} \delta_{ik} \right). \quad (5)$$

Throughout this work the coordinates x_i ($i = 0, 1, 2$) are referred to as (x, y, z) , with corresponding velocity components (u, v, w) . The equations are non-dimensionalized by reference velocity, density, and temperature $(U_{\text{ref}}^*, \rho_{\text{ref}}^*, T_{\text{ref}}^*)$. Prandtl number and the ratio of specific heat capacity for air are taken to be $Pr = 0.71$ and $\gamma = 1.4$ respectively. The $M_{\text{ref}} = 0.1$ subsonic Taylor-Green vortex (TGV) test case uses constant viscosity, whereas the compressible TGV use

dynamic viscosity $\mu(T)$ computed by Sutherland's law from [12] as

$$\mu(T) = \frac{1.4042T^{1.5}}{T + 0.40417}. \quad (6)$$

For an ideal Newtonian fluid, pressure can be calculated through the equation of state such that

$$p = (\gamma - 1) \left(\rho E - \frac{1}{2} \rho u_i u_i \right) = \frac{1}{\gamma M_{\text{ref}}^2} \rho T. \quad (7)$$

The heat flux and viscous terms are computed with standard 4th order central differencing schemes. For a constant grid spacing of Δx , the formula for first and second derivatives are given by

$$f'_i = \frac{f_{i-2} - 8f_{i-1} + 8f_{i+1} - f_{i+2}}{12\Delta x}, \quad (8)$$

$$f''_i = \frac{-f_{i-2} + 16f_{i-1} - 30f_i + 16f_{i+1} - f_{i+2}}{12\Delta x^2}. \quad (9)$$

For central difference schemes, two applications of a first derivative formula are not numerically equivalent to a single application of a second derivative. In the context of the Navier-Stokes equations, this can lead to undesirable odd-even decoupling phenomena [13]. To improve the robustness of the code, the viscous terms occurring in the momentum equations (2)

$$\frac{\partial \tau_{ij}}{\partial x_j} = \frac{\partial}{\partial x_j} \left(\mu \left[\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right] - \frac{2}{3} \mu \left[\frac{\partial u_k}{\partial x_k} \right] \delta_{ij} \right), \quad (10)$$

are rewritten in Laplacian form

$$\frac{\partial \tau_{ij}}{\partial x_j} = \mu \frac{\partial^2 u_i}{\partial x_j \partial x_j} + \frac{\mu}{3} \frac{\partial}{\partial x_i} \left(\frac{\partial u_k}{\partial x_k} \right) + \frac{\partial \mu}{\partial x_j} \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) - \frac{2}{3} \frac{\partial \mu}{\partial x_i} \frac{\partial u_j}{\partial x_j}, \quad (11)$$

where the first term is computed directly by the second derivative formula in equation (9).

Large-scale direct numerical simulation (DNS) of the compressible Navier-Stokes equations has considerable memory requirements, prompting the need for minimal-storage time-stepping algorithms. A low-storage explicit third-order Runge-Kutta time-stepping scheme [14] is used for time integration for all simulations in this work.

B. OpenSBLI high-order finite-difference code

All of the simulations in this work were carried out using the automatic source-code generation framework OpenSBLI [15], [16]. OpenSBLI is a Python-based finite-difference framework for structured meshes, that generates a C code in the Oxford Parallel Structured (OPS) embedded domain-specific-language (eDSL) [17]. Codes are generated to solve the equations specified by the user, using a selection of available numerical schemes. The base code is translated

by the OPS library to a range of common parallel programming paradigms. At present these include OpenMP, MPI, OpenMP+MPI for CPU/accelerators, and CUDA/OpenCL/OpenACC for GPUs. The results presented in this paper were performed using the MPI back-end for CPUs and the CUDA+MPI back-end for multi-GPU clusters.

C. Flux reconstruction

A flux reconstruction is performed to compute the convective terms of the Navier-Stokes equations in (1)-(3). The flux reconstruction is outlined in this section in the context of WENO/TENO schemes. The following is simplified to a 1D advection equation but is easily extended to systems of equations where the procedure is applied for each component of the system in turn. For a simple 1D hyperbolic equation of the form

$$U_t + f(U)_x = 0, \quad (12)$$

the flux term $f(U)_x$ at each discrete grid point x_i can be approximated by computing a flux reconstruction over the half-node locations $[x_{i-\frac{1}{2}}, x_{i+\frac{1}{2}}]$ with grid spacing Δx such that

$$\frac{\partial U}{\partial t} + \frac{1}{\Delta x} (\hat{f}_{i+\frac{1}{2}} - \hat{f}_{i-\frac{1}{2}}) = 0. \quad (13)$$

The full numerical stencil for a WENO/TENO scheme is constructed from a convex combination of lower order candidate stencils, as shown for TENO in Fig 1. Non-oscillatory behaviour is obtained by measuring the smoothness in each candidate stencil, and assigning it a weighting ω_r in the final reconstruction. Fluxes at the half-node locations $f_{i+\frac{1}{2}}$ and $f_{i-\frac{1}{2}}$ are reconstructed as a weighted sum of essentially non-oscillatory (ENO) interpolations over a set of r candidate stencils

$$\hat{f}_{i+\frac{1}{2}} = \sum_{r=0}^{K-3} \omega_r \hat{f}_{i+\frac{1}{2}}^{(r)}, \quad (14)$$

where $\hat{f}_{i+\frac{1}{2}}^{(r)}$ are the classic ENO interpolations given in [18], and ω_r are the non-linear weights satisfying the conditions

$$\omega_r \geq 0, \quad \text{and} \quad \sum_{r=0}^{K-3} \omega_r = 1. \quad (15)$$

Flux splitting is applied by summing upwind/downwind biased reconstructions $\hat{f} = \hat{f}^+ + \hat{f}^-$, using the well known local Lax-Friedrichs splitting method

$$f^\pm(U) = \frac{1}{2} (f(U) \pm \alpha U) \quad (16)$$

for a locally evaluated wave-speed α over the full numerical stencil. The negative flux contribution is obtained by reflecting the stencils in Fig 1 about the half-node reconstruction point $x_{i+\frac{1}{2}}$. For all simulations in this paper, the

reconstruction is performed in characteristic space as this achieves sharper shock capturing with reduced oscillations [8]. The implementation of the characteristic decomposition was validated for the present OpenSBLI code in [15], based on the work of [18]. For systems of equations the wave-speed α in the flux splitting (16) is the characteristic wave speed evaluated locally over the stencil.

D. Weighted Essentially Non-Oscillatory (WENO) weights

Implementation of the WENO scheme in the OpenSBLI code used in this work was presented in [15] and is briefly summarised here. In the standard Jiang-Shu formulation [18] (denoted WENO-JS), the non-linear weights ω_r for the flux reconstruction (14) are given for a scheme of order $2k - 1$ by

$$\omega_r = \frac{\alpha_r}{\sum_{s=0}^{k-1} \alpha_s}, \quad r = 0, \dots, k - 1, \quad (17)$$

where α_r is equal to

$$\alpha_r = \frac{d_r}{(\epsilon + \beta_r)^2}. \quad (18)$$

The polynomial smoothness indicators β_r [19] detect discontinuities and ensure that stencils containing shocks are weighted close to zero in the final flux reconstruction. The optimal weights d_r are the linear weights for smooth regions of the flow from [18] and the small parameter ϵ is equal to 1×10^{-6} . Shortcomings of the original WENO-JS formulation were addressed by Borges et al [3] in the improved WENO-Z scheme. WENO-Z maintains higher order of accuracy near critical points and exhibits lower levels of numerical dissipation via a reformulation of the non-linear weights (18). In addition to the standard measures of smoothness β_r in each of the smaller candidate stencils, the WENO-Z scheme introduces a global smoothness indicator τ_K over the full numerical stencil. The weighting (17) is modified such that

$$\alpha_r = d_r \left(1 + \frac{\tau_K}{\beta_r + \epsilon} \right), \quad r = 0, \dots, k - 1, \quad (19)$$

where β_r and d_r are unchanged from the WENO-JS formulation and τ_K is a global smoothness measure for a K th order scheme. In the case of the 5th order WENO-5Z the global smoothness indicator is given by $\tau_5 = |\beta_0 - \beta_2|$, representing the absolute difference in the smoothness at opposing sides of the full stencil. For higher order WENO-Z schemes the global smoothness indicators are given in [4].

In this work the WENO schemes are used both on their own in the entirety of the domain, and as part of a hybrid central-WENO approach. Hybrid schemes pair a non-dissipative central differencing scheme to dissipative shock capturing methods via a sensor. For the hybrid scheme used in this work, the Ducros sensor [20] is evaluated at the start




Fig. 1 Staggered candidate stencils for TENO schemes, adapted from [5].

of each Runge-Kutta sub-step on the entire domain as

$$\Phi = \frac{(\nabla \cdot \mathbf{u})^2}{(\nabla \cdot \mathbf{u})^2 + (\omega)^2 + \epsilon}, \quad (20)$$

The value of the sensor $\Phi \in [0, 1]$ is read in to the computational kernels performing reconstruction of the spatial fluxes for each direction in turn. The spatial fluxes are reconstructed by a non-dissipative central scheme for $\Phi < 0.67$ in smooth regions of the flow, and with the WENO scheme elsewhere. The Ducros split-skew symmetric formulation [21] is used for the central reconstruction to improve the numerical stability for the non-dissipative part.

E. Targeted Essentially Non-Oscillatory (TENO) schemes

TENO schemes differ from WENO in three fundamental ways: the staggered ordering of candidate stencils as in Fig 1, the complete removal of candidate stencils deemed to be non-smooth via a discrete cut-off function, and modified non-linear weights optimized for low numerical dissipation. The stencil staggering is said to reduce the number of candidate stencils that are crossed by a discontinuity [5], helping to enforce the underlying principles of the base ENO method. In the standard WENO formulation each candidate stencil is given a non-linear weighting, with candidate stencils crossing discontinuities given a low, albeit non-zero, weight. TENO schemes abandon this approach in favour of a discrete cut-off function that discards candidate stencils completely from the flux reconstruction if they are deemed to be sufficiently non-smooth. Candidate stencils that are considered smooth are included in the reconstruction with their ideal linear weight to further reduce numerical dissipation.

In reference to Fig 1, the 5th order reconstruction uses candidate stencils $\{S_0, S_1, S_2\}$, and the 6th order scheme uses $\{S_0, S_1, S_2, S_3\}$. The non-linear weights of [18] are reformulated with ideal weights d_r for a scheme of order K to be

$$\omega_r = \frac{d_r \delta_r}{\sum_{r=0}^{K-3} d_r \delta_r}, \quad (21)$$

where δ_r is a discrete cut-off function of the form

$$\delta_r = \begin{cases} 0 & \text{if } \chi_r < C_T \\ 1 & \text{otherwise} \end{cases} \quad (22)$$

for a tunable cut-off parameter C_T . The smoothness measures χ_r follow the same form as the weight normalization process in WENO

$$\chi_r = \frac{\gamma_r}{\sum_{r=0}^{K-3} \gamma_r}, \quad (23)$$

comprised of the WENO-Z [3] inspired form of non-linear TENO weights [5] given by

$$\gamma_r = \left(C + \frac{\tau_K}{\beta_r + \epsilon} \right)^q, \quad r = 0, \dots, K-3. \quad (24)$$

Polynomial smoothness indicators β_r are unchanged from the standard Jiang-Shu formulation [19], and are given explicitly for the TENO stencils of Fig 1 in [5]; the small parameter $\epsilon \sim 10^{-16}$ is used to avoid division by zero. A global smoothness indicator τ_K measures smoothness over the full numerical stencil, and is given for 5th and 6th order TENO schemes as $|\beta_0 - \beta_2|$ and $|\beta_3 - \frac{1}{6}(\beta_0 + \beta_2 + 4\beta_1)|$ respectively.

One feature of the TENO schemes is the optimization of the parameters C and q within the non-linear weights, and the tunable cut-off parameter C_T found in the discrete weight selection function. The parameters C and q control the levels of dissipation invoked by the non-linear weight, and are based on physically-motivated scale-separation mechanisms [5]. TENO schemes take the values of $q = 6$ and $C = 1$, which was shown to significantly enhance the discontinuity detection capability of the scheme [5]. C_T meanwhile is the parameter that determines whether a given candidate stencil is rejected or contributes to the flux reconstruction. Lower values of C_T are suitable for compressible turbulence where minimal numerical dissipation is desirable, but this comes at the cost of increased spurious oscillations around shock-waves. In the work of [5] the recommended optimal values of C_T for 5th and 6th order TENO schemes were stated to be 1×10^{-5} and 1×10^{-7} respectively.

F. Adaptive Targeted Essentially Non-Oscillatory (TENO-A) schemes.

A recent development of the base TENO schemes was proposed in [7], which will be referred to as TENO-A throughout. The TENO-A scheme applies adaptive control of the scheme's numerical dissipation in different regions of the flow. This is achieved by adding a shock sensor to modify the threshold value C_T defined in section II.E. In smooth regions of the flow fewer stencils will be discarded from the reconstruction when compared to the base TENO schemes. Near shocks the value of C_T is increased to robustly capture shocks and minimize spurious oscillations. In this work, the Ducros sensor [20] defined in equation (20) is used to enable adaptive dissipation control of the TENO-A scheme.

Alternatively, other sensors based on density or pressure gradients are also suitable to control the dissipation of the scheme [7]. The sensor is evaluated over the entire domain at the start of each Runge-Kutta sub-stage, to be used in determining an appropriate local value of C_T . The maximum value of the sensor over the local TENO stencil is used for the calculation of C_T . Mapping of the shock sensor Φ to the value of C_T is the same as that proposed by [7]

$$\begin{cases} g(\Phi) = (1 - \Phi)^4(1 + 4\Phi) \\ \bar{\beta} = \alpha_1 - \alpha_2(1 - g(\Phi)) \\ C_T = 10^{-\lfloor \bar{\beta} \rfloor} \end{cases} \quad (25)$$

for a Gauss bracket $\lfloor \bar{\beta} \rfloor$. The parameters α_1, α_2 are user defined based on the problem and are taken as $\alpha_1 = 10.5$ and $\alpha_2 = 5.5$ in this work. The resulting stencil cut-off parameter varies from $C_T = 10^{-10}$ in smooth regions to $C_T = 10^{-5}$ around shocks.

III. Problem specification

The Taylor-Green Vortex (TGV) is a well-known test case consisting of the laminar-turbulent transition of a decaying vortex. A seminal investigation of this problem was presented in the work of [9], who investigated both the inviscid and viscous subsonic TGV cases. Part of this study highlighted the intrinsic symmetries of the initial vortex condition. Despite the several decades that have passed since this work, the Taylor-Green vortex remains a physically interesting problem for the study of vortex dynamics and breakdown to turbulence [12, 22]. The TGV case is also widely used for assessing the efficiency of numerical schemes and often acts as a validation of CFD codes [11, 16, 23]. The standard subsonic ($M_{\text{ref}} = 0.1$, $Re = 1600$) case [10] is extended in this work up to $M_{\text{ref}} = 1.25$, where the presence of shocklets and compressibility effects become an important factor. The aim of this section is both to elucidate the intricate flow structures for a supersonic TGV, and to use it as a challenging test case for assessing shock-capturing schemes. A comparable physical study of high Mach number TGV was presented by [12] up to Mach 2, albeit at a lower Reynolds number of $Re = 400$. At this Reynolds number the range of turbulent scales is not large enough to show significant differences between the numerical schemes.

The initial condition is taken from [10], solving the compressible Navier-Stokes equations on a domain of size $0 \leq x \leq 2\pi L$, $0 \leq y \leq 2\pi L$, and $0 \leq z \leq 2\pi L$ as in Fig 2. Periodic boundary conditions were applied in all directions. The domain was initialised as

$$u(x, y, z, t = 0) = \sin\left(\frac{x}{L}\right) \cos\left(\frac{y}{L}\right) \cos\left(\frac{z}{L}\right), \quad (26)$$

$$v(x, y, z, t = 0) = -\cos\left(\frac{x}{L}\right) \sin\left(\frac{y}{L}\right) \cos\left(\frac{z}{L}\right), \quad (27)$$

$$w(x, y, z, t = 0) = 0, \quad (28)$$

$$p(x, y, z, t = 0) = \frac{1}{\gamma M_{\text{ref}}^2} + \frac{1}{16} \left(\cos\left(\frac{2x}{L}\right) + \cos\left(\frac{2y}{L}\right) \right) \left(2 + \cos\left(\frac{2z}{L}\right) \right). \quad (29)$$

Reynolds, and Prandtl numbers were set to $\text{Re} = 1600$ and $\text{Pr} = 0.71$ respectively, the ratio of heat capacities was taken to be $\gamma = 1.4$. The length L is taken to be 1, with the initial density evaluated from the equation of state (7) with $T(x, y, z, t = 0) = 1$. This deterministic initial condition triggers large-eddies that break down to turbulence and, as we shall see, create eddy shock-waves at high Mach numbers.

A fixed non-dimensional time-step of $\Delta t = 5 \times 10^{-4}$ is set throughout, based on the stability requirements for the finest $N = 1024^3$ mesh simulation. Larger time-steps could be used to improve computational efficiency for the coarser meshes. The TGV case is advanced until a non-dimensional time of $t = 20$. To assess the numerical schemes, comparison is made for the kinetic energy integrated over the domain

$$E_k = \frac{1}{\rho_{\text{ref}} \Omega} \int_{\Omega} \frac{1}{2} \rho u_j u_j d\Omega, \quad (30)$$

and total viscous dissipation rate

$$\varepsilon^T = \varepsilon^S + \varepsilon^D = \frac{1}{\rho_{\text{ref}} \text{Re}} \int_{\Omega} \mu \left(\epsilon_{ijk} \frac{\partial u_k}{\partial x_j} \right)^2 d\Omega + \frac{4}{3\rho_{\text{ref}} \text{Re}} \int_{\Omega} \mu \left(\frac{\partial u_j}{\partial x_j} \right)^2 d\Omega, \quad (31)$$

corresponding to the solenoidal (enstrophy) and dilatational contributions to the viscous dissipation rate for compressible turbulence [24, 25]. Q-criterion is calculated in this work from [26] as

$$Q = \frac{1}{2} \left[\left(\frac{\partial u_i}{\partial x_i} \right)^2 - \frac{\partial u_i}{\partial x_j} \frac{\partial u_j}{\partial x_i} \right]. \quad (32)$$

IV. Results

A. Sensitivity to grid refinement

Fig 3 shows the compressible Taylor-Green vortex problem at Mach 1 for meshes ranging from 0.25 million to 1 billion grid points, computed with the 6th order TENO scheme. The kinetic energy in Fig 3 (a) is largely grid independent, with only the coarsest 64^3 case showing a clear deviation from the converged result. At 64^3 the kinetic energy is underestimated between $5 < t < 15$ and over-estimated at $t > 15$. The total dissipation (31) is displayed in Fig 3 (b) and shows a much greater dependency on mesh refinement. Meshes between 64^3 and 256^3 substantially underestimate the peak dissipation rate. Above 512^3 a converged solution is obtained. Comparison of Mach number effects in the next section use the 512^3 grid, while the scheme comparisons are performed on coarser 128^3 and 256^3




Fig. 2 Initial condition for the 3D Taylor-Green vortex test case. The domain is a periodic box of size $2\pi^3$. A Q-criterion of 0.25 is displayed, coloured by the u component of velocity.

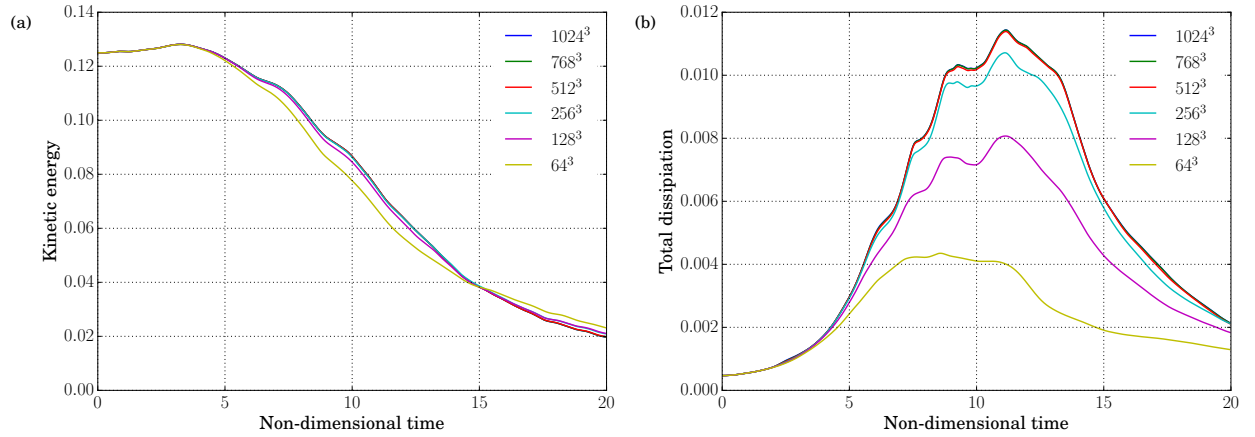


Fig. 3 Grid sensitivity of the $M_{\text{ref}} = 1$ compressible Taylor-Green vortex problem for (a) kinetic energy and (b) total dissipation with the TENO6 scheme.

meshes to assess how the schemes perform as ILES when limited by grid resolution. The kinetic energy and viscous dissipation rate were calculated in the domain every 100 time-steps and integrated over the volume.

B. Physical effects of increasing Mach number

Fig 4 shows the effect of Mach number scaling for the TGV problem over the range $M_{\text{ref}} = [0.1, 1.25]$. In Fig 4 (a) the first significant deviation from the subsonic $M_{\text{ref}} = 0.1$ case is observed at $M_{\text{ref}} = 0.5$ in the kinetic energy distribution. Unlike the low-subsonic case, the compressible problem has regions of increasing kinetic energy rather than a constant decline. This is consistent with the compressible TGV simulations of [12], where increases in the kinetic energy were reported during the time interval $2 < t < 4$. In compressible flows the average pressure-work can exceed the viscous dissipation rate, as internal energy is converted into kinetic energy during the early stages of the problem.




Fig. 4 The effect of increasing Mach number for the compressible $Re = 1600$ Taylor-Green vortex problem on 512^3 grids, displaying (a) kinetic energy and (b) total dissipation.

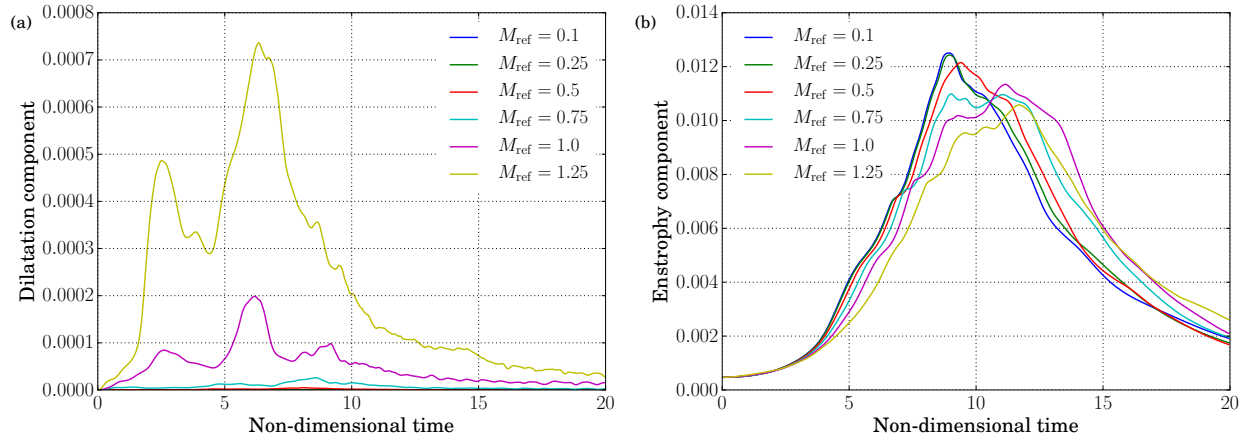


Fig. 5 The effect of increasing Mach number for the compressible $Re = 1600$ Taylor-Green vortex problem on 512^3 grids, displaying (a) dilatational dissipation and (b) solenoidal dissipation.

The increase in kinetic energy is most obvious at $M_{\text{ref}} = 1.25$, but is clearly evident once the reference Mach number reaches 0.75. There is good agreement between the $M_{\text{ref}} = 0.1$ and $M_{\text{ref}} = 0.25$ results, suggesting that incompressible approximations hold at $M_{\text{ref}} = 0.25$.

Fig 4 (b) shows the total viscous dissipation for the same range of Mach numbers. A large spread of dissipation curves is observed, the main trends being a delay and a flattening of the peak dissipation rate for increasing Mach numbers. Similar to the kinetic energy, peak dissipation rate for the $M_{\text{ref}} = 0.25$ curve matches closely to the $M_{\text{ref}} = 0.1$ result. For the higher Mach numbers the dissipation rate is significantly higher in the $12 < t < 20$ region than in the incompressible limit. At Mach 1.25 an additional peak is seen for the early $2 < t < 4$ time period, similar to the one observed in [12]. The total viscous dissipation can be decomposed into the dilatational and solenoidal components shown in Fig 5 (a) and (b) respectively. There are two distinct peaks in the dilatational contribution for the Mach 1 and Mach 1.25 curves, the latter corresponding to the peak seen in Fig 4 (b) for $2 < t < 4$. Shocklets are generated

during the early stages of the high Mach number simulations, a compressibility effect that is not present at lower Mach numbers. At Mach 1.25 the magnitude of the dilatation is still small compared to the solenoidal contribution, but at even higher Mach numbers is expected to match the solenoidal dissipation [12]. The combination of vortex dynamics, shocklets, and small-scale turbulence makes this a challenging problem for numerical schemes.




Fig. 6 Formation of shock-waves during the early stages of the compressible $M_{\text{ref}} = 1.25$ Taylor-Green vortex case. Showing (left) dilatation rate and (right) local Mach number at a simulation time of $t = 2$.




Fig. 7 Formation of shock-waves in the compressible Mach 1.25 Taylor-Green vortex case. Showing (a) discontinuous Mach profile along the $y, z = \pi$ line and (b) time evolution of the maximum Mach number in the domain.

To further demonstrate the physical complexity of this test case, Fig 6 shows a three-dimensional surface during the initial stages of the Mach 1.25 simulation on the $N = 512^3$ mesh. The simulation time is taken to be $t = 2$, corresponding to the first dilatational peak shown in Fig 5. The left and right images show the instantaneous dilatation rate $(\nabla \cdot \mathbf{u})$, and

the local Mach number based on the velocity magnitude and speed of sound

$$M = \frac{|U|}{a} = \frac{\sqrt{u^2 + v^2 + w^2}}{a}. \quad (33)$$

Recalling the observation that negative dilatation rates are associated with shock-waves [2], we see multiple thin sharply-captured shock-waves throughout the domain in white and blue. Depending on the orientation, these shock-waves propagate both inwards and outwards relative to the periodic boundaries of the domain. The corresponding Mach number surface in Fig 6 (right) shows the high-speed regions adjacent to the shock-waves. At this simulation time of $t = 2$, regions of the flow have locally accelerated from $M_{\text{ref}} = 1.25$ up to Mach 1.5. On the upper surface of the Mach number plot five distinct low-speed regions can be seen. As we will see, the four of these outside of the centre correspond to counter-rotating flow driven by the relative propagation of the initial shock-waves.

Fig 7 (a) shows two line profiles in x on the centreline plane of $y, z = \pi$. The smooth black line is the initial condition (26) evaluated at $t = 0$. The blue line corresponds to the simulation at $t = 2.5$, to demonstrate the discontinuous shocked regions of the flow. Either side of the centreline, a sharp well-captured shock-wave is observed as a sudden drop in the Mach number from $M = 1.5$ to $M = 0.45$. The smooth initial distribution rapidly develops into a discontinuous shock-wave due to the non-linearity of the compressible Navier-Stokes equations. For subsequent time snapshots, the rounded top of the Mach profile sharpens and the shock-waves propagate away from each other. In other parts of the domain shock-waves cross the periodic boundaries and interfere with each other. In Fig 7 (b), the time evolution of the maximum Mach number within the volume is shown. From the initial reference Mach number of 1.25, a peak Mach number of 2 is observed at $t = 6$. Locally supersonic flow is observed even in the latter stages of the simulation, up until a time of $t = 14$.

Turning the attention to the small-scale structures for this transitional case, Fig 8 shows the generation of vortical structures at a Q -criterion of $Q = 2.5$, coloured by the u component of velocity. For the plots clockwise from the upper left corner, the flow is shown at simulation times of $t = [2, 7, 15, 20]$. At $t = 2$ the propagation of the shock-waves observed in Fig 6 (left) can be seen to be generating non-zero values of Q . The opposite colouring of the shock-surfaces demonstrates the relative motion of the shock-waves in the early stages of the simulation. The second Q plot at $t = 7$ corresponds to the second peak of dilatational dissipation from Fig 5. Two distinct regions of vorticity generation are observed for this Q value at $z = \frac{\pi}{2}$ and $z = \frac{3\pi}{2}$. The relative motion of the fluid has led to the creation of thick vortical structures that merge, twist, and reconnect with one another.

Fig 9 shows a top-down view of the same structures at $t = 7$, albeit with a larger value of $Q = 12.5$ to highlight the largest structures. In each of the two structures from Fig 8, the relative motion of the initial shock-waves is seen to generate elongated vortex structures along the domain boundaries and the principal axes $x, y = \pi$. The shock-waves propagating along these axes pass through the periodic boundaries and reflect from one another. This can be seen from

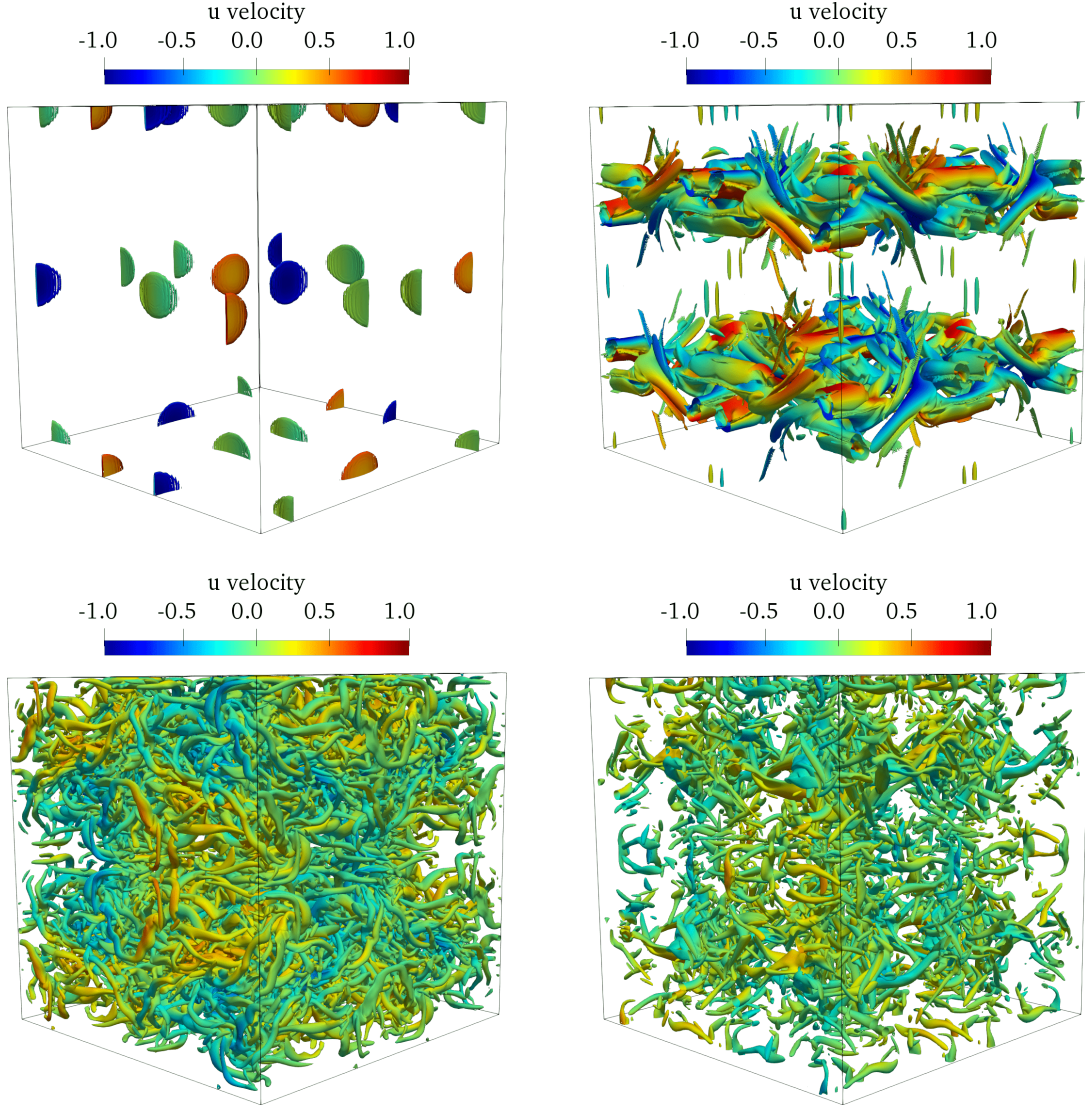


Fig. 8 Vortical structures in the $M_{\text{ref}} = 1.25$ Taylor-Green vortex case at $Q = 2.5$ coloured by the x component of velocity. Showing simulation times of $t = [2, 7, 15, 20]$ in the clockwise direction starting from the upper left corner.

the opposing regions of dark red and blue in close proximity to one another. Four distinct quadrants are formed due to the symmetries of the initial condition 26. On the perimeters of each of these quadrants, small vortex structures are generated from the relative motion of the fluid along the principal axes. We refer once more to the upper Mach number surface of Fig 6, where the tails of the rotating low-Mach regions match the location of these small-scale vortices. As this simulation time corresponds to the largest peak in dilatational dissipation from Fig 5, we conclude that the reinforcement and reflections of the propagating shock-waves play a crucial role in determining the viscous dissipation. As such, robust shock-capturing schemes that can sharply capture the shock-waves are essential for the correct prediction of viscous dissipation rates.




Fig. 9 Top view of the vortical structures in the $M_{\text{ref}} = 1.25$ Taylor-Green vortex case at $Q = 12.5$ coloured by velocity. Showing an $(x-y)$ view of the simulation at a time of $t = 7$.

The final two Q -criterion plots in Fig 8 show the breakdown to turbulence in the latter stages of the simulation. Large vortices are observed to populate the entire computational box as the turbulent mixing rates increase. At the final $t = 20$ stage of the simulation, the breakdown and cascade to smaller turbulent scales is evident. The four-quadrant structure in the $(x-y)$ plane is shown at $z = \pi$ for the same four simulation times of $t = [2, 7, 15, 20]$ in Fig 10. The quantity plotted is the instantaneous magnitude of density gradients

$$|\nabla \rho| = \sqrt{\left(\frac{\partial \rho}{\partial x}\right)^2 + \left(\frac{\partial \rho}{\partial y}\right)^2 + \left(\frac{\partial \rho}{\partial z}\right)^2}. \quad (34)$$

Unlike the subsonic case which has only minor density variations, the supersonic TGV shows numerous regions of large density gradients. Due to the trigonometric symmetries of the problem (26), eight distinct shock-waves are observed in Fig 10 (a) at $t = 2$ on this centreline plane. The shocks propagate away from the central $(x, y = \pi)$ axis point. By $t = 7$ the shock-waves have reflected and formed elongated structures along the two axes, which were observed previously in Fig 9. The four quadrants are qualitatively similar to the compressible TGV simulations shown in Fig 4 of [12]. At $t = 15$ where the flow has become subsonic (Fig 7), small regions of high-density gradients are still observed to propagate along the axes and domain boundaries. The relative motion of the fluid along the axes generates counter-rotating flow in each of the four visible quadrants. Fig 10 (d) shows the structures present at the final simulation time of $t = 20$. Despite the breakdown to small-scale structures observed in Fig 8, at the larger scales the four quadrant symmetry is still present during this late stage of the simulation.




Fig. 10 Magnitude of density gradients $|\nabla\rho|$ for the TGV at $M_{\text{ref}} = 1.25$ on a 512^3 grid. Showing the flow structures at simulation times of $t = [2, 7, 15, 20]$ in (a) to (d), on the centreline (x - y) plane located at $z = \pi$.

Figures 11 (a)-(d) show the magnitude of density gradients in the (x - z) plane on the $y = \pi$ centreline, during the early-to-mid stages of the simulation ($t \in [0, 12.5]$). In this plane at $t = 2.5$ well-defined normal shocks are observed, straighter than those seen in the (x - y) plane of Fig 10 (a). The middle two shock-waves travel in opposite directions to the ones at the domain boundaries. By $t = 5$ in Fig 10 (b), the initial shock-waves have interacted with one another, with regions of both constructive and destructive interference. At this time multiple shocks can be seen moving obliquely relative to the coordinate axes. The large-scale shock-structures persist at $t = 7.5$, shortly after the peak local Mach number shown in Fig 7 (b). Finally at $t = 7.5$ and $t = 10$, two layers of vortical structures are observed: one in the region of $0 < z < \pi$, and the other within $\pi < z < 2\pi$. These large-scale vortex structures correspond to the two layers of Q surfaces seen in Fig 8 and Fig 9. Finally at $t = 12.5$ in Fig 11 (e), eight compact high density-gradient regions are observed in the cross-section of the vortex structures. The travelling shock-waves are interconnected and interfere with these vortex tubes. As a result, correctly capturing the shock-waves is essential to reproduce the vortex dynamics and

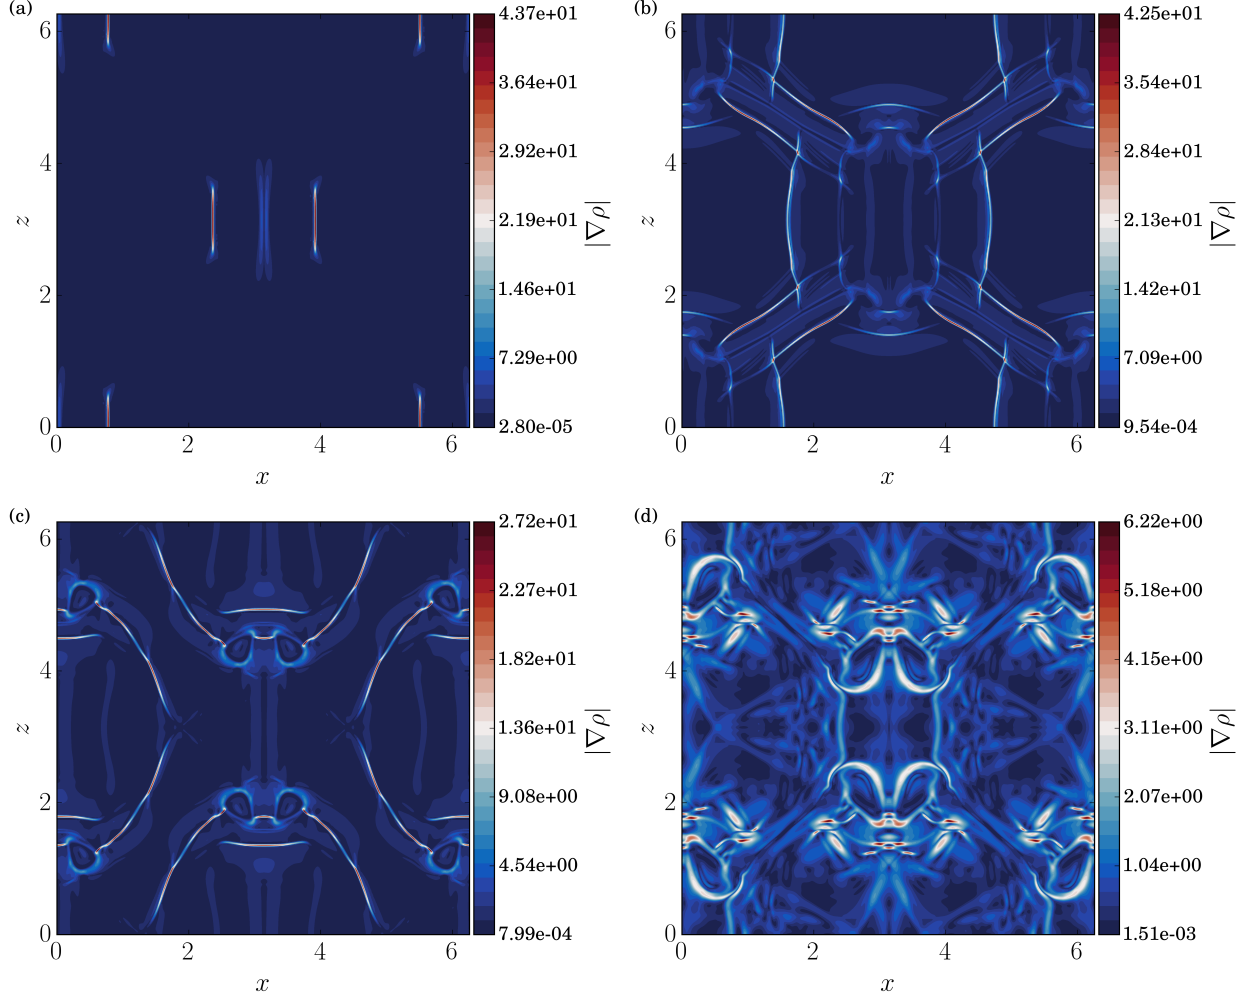


Fig. 11 Magnitude of density gradients $|\nabla\rho|$ for the TGV at $M_{\text{ref}} = 1.25$ on a 512^3 grid. Showing the flow structures at simulation times of $t = [2.5, 5, 7.5, 10]$ in (a) to (d), on the centreline (x - z) plane located at $y = \pi$.

subsequent breakdown to turbulence in the latter stages of the problem.

In this section we have seen the additional complexity that is present in a supersonic version of the Taylor-Green vortex test case. In contrast to the standard subsonic $M_{\text{ref}} = 0.1$ case, the kinetic energy distribution is no longer found to be monotonically decreasing (Fig 4 (a)). In the early stages of the simulation, internal energy is converted into kinetic energy which can exceed the viscous dissipation rate. While the incompressible approximation is still largely valid at $M_{\text{ref}} = 0.25$ (Fig 4 (b)), compressibility effects quickly become important at higher Mach numbers and lead to a delay in the peak dissipation rate. It is possible that this is due to the stabilising effect of increasing Mach number on the growth rate of instabilities, seen previously in compressible mixing-layers and turbulent spots [27, 28]. Temporally decaying shock-waves were observed throughout the domain (Fig 6), demonstrating the need for shock-capturing capabilities for this test case. Despite the similarities to the well-studied subsonic case, the demonstrated compressibility effects make the supersonic Taylor-Green vortex an interesting physical problem in its own right. For the purposes of this work,

the supersonic case poses a suitably challenging problem to test the performance of shock-capturing schemes in the presence of compressible turbulence.

C. Shock-capturing scheme comparison




Fig. 12 Taylor-Green vortex $M_{\text{ref}} = 0.1$ scheme comparison on coarse 256^3 grids for (a) kinetic energy and (b) enstrophy compared to a reference fine 512^3 spectral solution (blue).

In this section, results are shown from simulations at $M_{\text{ref}} = 0.1$ and $M_{\text{ref}} = 1.25$, using the 5th order WENO-JS, 5th and 7th order WENO-Z, 5th and 6th order TENO schemes, 6th order TENO6-A and a hybrid central-WENO-Z scheme. In the $M_{\text{ref}} = 0.1$ subsonic case the schemes are compared for 256^3 grids compared to a 512^3 spectral result from [29]. In the incompressible limit shock capturing schemes are not required for this problem, but the existence of well validated reference solutions allows a comparison of the small-scale resolving ability of the schemes.

Fig 12 (a) shows that for $M_{\text{ref}} = 0.1$, the kinetic energy is largely independent of the scheme at 256^3 ; only the WENO schemes slightly underestimate the reference result. The excessive numerical dissipation of shock capturing schemes is clear to see, however, in Fig 12 (b), where all of the schemes severely underestimate the dissipation rate compared to a non-dissipative 4th order central scheme. The hybrid scheme (not shown) is identical to the central result for the subsonic $M_{\text{ref}} = 0.1$ case as the shock sensor is never activated for this problem. The standard WENO5-JS is the worst performing scheme, underestimating the dissipation rate by up to 40%. The WENO-Z results are better, especially when considering the minor modifications the scheme requires compared to the original WENO-JS formulation. The addition of the global smoothness measure to the non-linear weights (19) substantially lowers the numerical dissipation of the WENO scheme, making it difficult to recommend the standard WENO-JS over WENO-Z when numerical dissipation is an important consideration. While still below the non-dissipative central scheme, both the 5th and 6th order TENO schemes perform well; low numerical dissipation is achieved while the schemes retain a good shock capturing ability.

Having demonstrated the numerical dissipation inherent to each of the schemes for a low Mach number problem, the natural extension is to see how they fare at higher Mach numbers where shock capturing is also required. Figures 13 (a)




Fig. 13 Taylor-Green vortex $M_{\text{ref}} = 1.25$ scheme comparison on coarse 256^3 grids for (a) kinetic energy and (b) total dissipation compared to a fine 768^3 mesh TENO6 solution (dashed line).

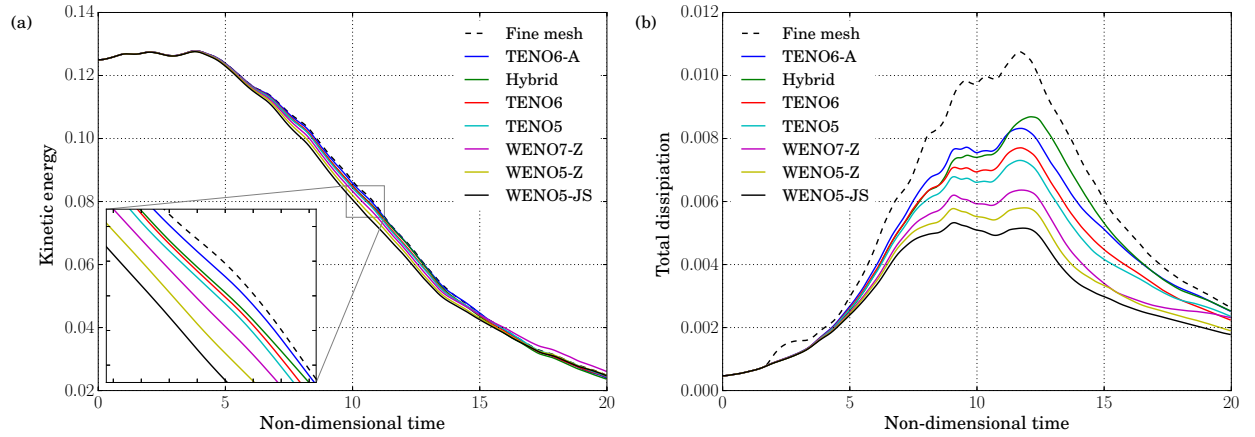


Fig. 14 Taylor-Green vortex $M_{\text{ref}} = 1.25$ scheme comparison on coarse 128^3 grids for (a) kinetic energy and (b) total dissipation compared to a fine 768^3 mesh TENO6 solution (dashed line).

and (b) show the kinetic energy and total dissipation at Mach 1.25 compared to the 768^3 fine mesh result from Fig 3. All of the schemes agree well for the kinetic energy evolution, but a large variance is again seen for the dissipation rate. The WENO5-JS scheme is the worst performer again, but the spread between the schemes is not as large as in the subsonic case. The TENO schemes prove to be less dissipative than WENO also when shocks are present, but fall short of the hybrid central-WENO-Z scheme. The hybrid and adaptive TENO-6A schemes with the dilatation/vorticity-based shock sensor (20) perform well and give the closest agreement to the fine mesh solution, but may not be suitable for every flow type [2]. We also observe that on the 256^3 mesh, all of the schemes under-predict the dilatational contribution at $t = 2.5$, which was shown for $M_{\text{ref}} = 1.25$ in fig 5 (a). While the standard TENO schemes are slightly more dissipative than the hybrid and adaptive TENO-6A methods, they rely only on one tunable parameter and avoid potential tuning of shock sensors for different flow configurations.

Figures 14 (a) and (b) show the scheme comparison on a coarser $N = 128^3$ mesh. The hierarchy of scheme

performance follows the same ordering as before, with the hybrid and TENO-6A schemes both performing well. At $N = 128^3$ there is a wider spread of the kinetic energy distribution in figure 14 (b) however, with the more dissipative schemes under-predicting the kinetic energy. This is most apparent for $7 < t < 13$, where the peak dissipation occurs during the early stages of the transition. Interestingly, the hybrid scheme under-performs the TENO-6A scheme in the first half of the simulation, but fares better for $t > 12$, as the flow breaks down to smaller scales.

D. Computational cost comparison

Table 1 Runtime for 1000 iterations of the TGV problem at Mach 1.25 on a 256^3 grid. CPU: Intel Skylake node (40 cores @ 2 GHz, 40 MPI, Intel 17.0 -O3 -fp-model fast). GPU: 1x NVIDIA Volta 16GB V100 (CUDA 9.0, nvcc -O3). The relative computational cost is compared to the WENO5-JS scheme.

Scheme	CPU Runtime (s)	Cost	GPU Runtime (s)	Cost
WENO5-JS	1202.3	-	190.8	-
WENO5-Z	1258.6	1.05	195.9	1.03
WENO7-Z	1651.8	1.37	236.7	1.24
TENO5	1377.1	1.15	201.8	1.06
TENO6	1683.8	1.40	216.5	1.13
TENO6-Adaptive	1772.3	1.47	223.7	1.17
Hybrid Central4-WENO5-Z	1016.8	0.84	179.0	0.94

In addition to the dissipation and shock-resolving ability of the schemes, an important consideration for DNS is the computational cost they incur. Table 1 shows the runtime for 1000 iterations of the TGV problem on the 256^3 grids used for the scheme comparison in figure 13. In each case multiple runs were performed to determine if there was any variability in the runtime, but repeated runs of the code were always consistent to one decimal place. Each scheme was tested for both a CPU and GPU-based platform, the domain specific language and code-generation allows for the same single source-code to be translated to a wide range of architectures. For both platforms the relative runtime was compared to a standard WENO5-JS scheme, to see the increase in computational cost required to achieve low-dissipative shock capturing. Performance of the OPS library used for parallel execution of the code has been reported previously in [30, 31], for large-scale DNS on multiple architectures. Additionally, parallel efficiency of the TENO shock-capturing schemes has been reported in [32], for weak- and strong-scaling on multi-GPU configurations.

WENO5-Z is an attractive option if very low dissipation is not the primary concern, it performs well for only a 3 – 5% increase in computational cost over WENO5-JS. It is difficult to recommend the use of WENO-JS for this type of problem considering the improvements made by WENO-Z for a negligible increase in cost. The 7th order WENO-Z scheme has both an additional stencil to compute and a 2-point wider numerical stencil than the 5th order scheme, making it up to 37% more expensive in these tests. A 5th order TENO scheme proved marginally more expensive than WENO5-JS, owing to the need to evaluate the discrete cut-off functions in equation 22. The 5th and 6th order TENO

schemes were 15% and 40% more expensive on CPU and only 6% and 13% more expensive on GPU.

There are several possible explanations for these architecture differences, including the effect of numerical stencil width, and how well a given architecture optimizes the branching ‘if-else’ conditions found in the TENO and hybrid schemes. In terms of the numerical stencil, the staggered-stencil TENO6 approach [5, 6] has the same number of candidate stencils and weights to compute as WENO7-Z, but has a narrower stencil of $[i - 2, i + 3]$, compared to $[i - 3, i + 4]$ for WENO7-Z (for a combined upwind/downwind reconstruction). This reduces the amount of data that has to be fetched to the computational kernel at each grid point, which seems to be an important factor for our GPU code. Additionally, the numerical stencil width determines the depth of MPI halo-exchanges required, which is an important consideration for much larger distributed memory simulations [31].

A second performance consideration is the effect of branch divergence, which is often a cause for concern for computational performance on GPU platforms. Despite this concern, the branching ‘if-else’ conditions in the TENO schemes seem to have better relative performance on the GPU than CPU. An alternative approach would be to compute the discrete TENO function (22) as the difference between the weight and threshold value, with a ceiling/floor function applied to obtain the binary stencil selection. The adaptive TENO-6A scheme has a further small increase in cost owing to the need to evaluate the shock sensor. Despite concerns about load balancing for hybrid methods that switch between two numerical schemes of differing cost, the hybrid method performs well and is cheaper than the WENO5-JS scheme.

V. Conclusion

A selection of low-dissipative shock capturing schemes from the ENO family have been assessed in the context of a compressible Taylor-Green vortex problem up to supersonic Mach numbers, with a Reynolds number of $Re = 1600$. At this Reynolds number a wide range of scales are present to stress the numerical schemes. Extension of the classic incompressible TGV to Mach 1.25 highlighted the formation of multiple strong shock-waves throughout the domain. These shock-waves were observed to interact and merge, resulting in complex and highly-unsteady shock-structures that drive the early stages of the laminar-turbulent transition of the decaying vortex. Unlike the widely-used low Mach number version of the TGV problem, the Mach 1.25 case appropriately stresses the shock-capturing schemes for shocked compressible turbulence. The challenging nature of the supersonic TGV places substantial demands on the numerical schemes to both robustly capture multiple moving shock-waves, and to resolve small-scale turbulence.

All of the schemes were able to reproduce similar shaped distributions for the kinetic energy and total dissipation rate. In contrast to the incompressible case, the TGV at higher Mach numbers was shown to have a kinetic energy distribution that was no longer monotonically decreasing in time. In the early stages of the simulation, internal energy is converted into kinetic energy that can exceed the viscous dissipation rate. At higher Mach numbers the dilatational contribution to the dissipation rate is no longer negligible. Mach number scaling of the TGV problem was investigated, with a significant delay in the peak dissipation rate with increasing Mach number. It is possible that the delay in the

peak dissipation rate is due to compressibility effects that damp the growth rates of instabilities, leading to a delay in the transition to turbulence. From an initial reference Mach number of 1.25, regions of the flow accelerated to a local peak of Mach 2, a third of the way into the simulation.

For ILES of the compressible TGV on coarser grids ($N = 128^3, 256^3$) at $M_{\text{ref}} = 1.25$, there was a clear and consistent ordering for the performance of the schemes compared to a fine mesh result. In general the low-dissipative and higher-order schemes performed better, with the TENO schemes shown to be consistently less-dissipative than their WENO counterparts. Conventional WENO schemes proved excessively dissipative for resolving the compressible turbulence unless paired with a hybrid central scheme. Low-dissipative TENO schemes are a significant improvement over WENO for comparable computational cost, making them better suited to DNS or ILES of compressible turbulence. In particular, the adaptive TENO6-A scheme of [7] with an in-built shock sensor performed well when resolving compressible turbulence on coarse grids. The standard TENO schemes use a global constant parameter for dissipation control, which is overly dissipative in smooth regions of the flow. The inclusion of a shock-sensor to enable local adaptation of this parameter is a simple way to improve the resolving performance of an existing TENO scheme in a code.

The best performing TENO6-A scheme was only 17% more expensive than a conventional WENO scheme despite offering substantial improvements in flow resolution. Despite concerns of potential load-balancing issues, the hybrid scheme was the computationally cheapest scheme tested, with comparable performance to TENO6-A on the coarse meshes. While low-dissipative high-order shock-capturing schemes can offer a simpler implementation compared to a coupled hybrid approach, hybrid schemes are still one of the most efficient options for simulating compressible fluid flows.

Acknowledgements

David J. Lusher was supported by an EPSRC Centre for Doctoral Training grant (EP/L015382/1). The work has been performed using resources provided by the ‘Cambridge Service for Data Driven Discovery’ (CSD3, <http://csd3.cam.ac.uk>) system operated by the University of Cambridge Research Computing Service (<http://www.hpc.cam.ac.uk>) funded by EPSRC Tier-2 capital grant EP/P020259/1. The authors would also like to acknowledge the use of the IRIDIS 5 High Performance Computing Facility, and associated support services at the University of Southampton. Data from this work will be made available on the University of Southampton institutional repository.

References

- [1] Pirozzoli, S., “Numerical Methods for High-Speed Flows,” *Annual Review of Fluid Mechanics*, Vol. 43, No. 1, 2011, pp. 163–194. <https://doi.org/10.1146/annurev-fluid-122109-160718>.
- [2] Johnsen, E., Larsson, J., Bhagatwala, A. V., Cabot, W. H., Moin, P., Olson, B. J., Rawat, P. S., Shankar, S. K., Sjögreen, B., Yee, H., Zhong, X., and Lele, S. K., “Assessment of high-resolution methods for numerical simulations of compressible turbulence with shock waves,” *Journal of Computational Physics*, Vol. 229, No. 4, 2010, pp. 1213 – 1237. <https://doi.org/https://doi.org/10.1016/j.jcp.2009.10.028>.
- [3] Borges, R., Carmona, M., Costa, B., and Don, W. S., “An improved weighted essentially non-oscillatory scheme for hyperbolic conservation laws,” *Journal of Computational Physics*, Vol. 227, No. 6, 2008, pp. 3191–3211. <https://doi.org/https://doi.org/10.1016/j.jcp.2007.11.038>.
- [4] Castro, M., Costa, B., and Don, W. S., “High order weighted essentially non-oscillatory WENO-Z schemes for hyperbolic conservation laws,” *Journal of Computational Physics*, Vol. 230, No. 5, 2011, pp. 1766–1792. <https://doi.org/https://doi.org/10.1016/j.jcp.2010.11.028>.
- [5] Fu, L., Hu, X. Y., and Adams, N. A., “A family of high-order targeted ENO schemes for compressible-fluid simulations,” *Journal of Computational Physics*, Vol. 305, 2016, pp. 333–359. <https://doi.org/https://doi.org/10.1016/j.jcp.2015.10.037>.
- [6] Fu, L., Hu, X. Y., and Adams, N. A., “Targeted ENO schemes with tailored resolution property for hyperbolic conservation laws,” *Journal of Computational Physics*, Vol. 349, 2017, pp. 97 – 121. <https://doi.org/https://doi.org/10.1016/j.jcp.2017.07.054>.
- [7] Fu, L., Hu, X. Y., and Adams, N. A., “Improved Five- and Six-Point Targeted Essentially Nonoscillatory Schemes with Adaptive Dissipation,” *AIAA Journal*, Vol. 57, No. 3, 2018, pp. 1143–1158. <https://doi.org/https://doi.org/10.2514/1.J057370>.
- [8] Brehm, C., Barad, M. F., Housman, J. A., and Kiris, C. C., “A comparison of higher-order finite-difference shock capturing schemes,” *Computers and Fluids*, Vol. 122, 2015, pp. 184–208. <https://doi.org/https://doi.org/10.1016/j.compfluid.2015.08.023>.
- [9] Brachet, M. E., Meiron, D. I., Orszag, S. A., Nickel, B. G., Morf, R. H., and Frisch, U., “Small-scale structure of the Taylor–Green vortex,” *Journal of Fluid Mechanics*, Vol. 130, 1983, pp. 411–452. <https://doi.org/https://doi.org/10.1017/S0022112083001159>.

- [10] DeBonis, J., “Solutions of the Taylor-Green Vortex Problem Using High-Resolution Explicit Finite Difference Methods,” *51st AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*, Aerospace Sciences Meetings, American Institute of Aeronautics and Astronautics, 2013. <https://doi.org/https://doi.org/10.2514/6.2013-382>.
- [11] Bull, J. R., and Jameson, A., “Simulation of the Taylor–Green Vortex Using High-Order Flux Reconstruction Schemes,” *AIAA Journal*, Vol. 53, No. 9, 2015, pp. 2750–2761. <https://doi.org/https://doi.org/10.2514/6.2014-3210>.
- [12] Peng, N., and Yang, Y., “Effects of the Mach number on the evolution of vortex-surface fields in compressible Taylor-Green flows,” *Phys. Rev. Fluids*, Vol. 3, 2018, p. 013401. <https://doi.org/https://doi.org/10.1103/PhysRevFluids.3.013401>.
- [13] Sandham, N., Li, Q., and Yee, H., “Entropy Splitting for High-Order Numerical Simulation of Compressible Turbulence,” *Journal of Computational Physics*, Vol. 178, No. 2, 2002, pp. 307 – 322. <https://doi.org/https://doi.org/10.1006/jcph.2002.7022>, URL <http://www.sciencedirect.com/science/article/pii/S0021999102970221>.
- [14] Carpenter, M. H., and Kennedy, C. A., “Fourth-order 2N-storage Runge-Kutta schemes,” *NASA Langley Research Center, NASA TM 109112.*, 1994.
- [15] Lusher, D. J., Jammy, S. P., and Sandham, N. D., “Shock-wave/boundary-layer interactions in the automatic source-code generation framework OpenSBLI,” *Computers & Fluids*, Vol. 173, 2018, pp. 17 – 21. <https://doi.org/https://doi.org/10.1016/j.compfluid.2018.03.081>.
- [16] Jacobs, C. T., Jammy, S. P., and Sandham, N. D., “OpenSBLI: A framework for the automated derivation and parallel execution of finite difference solvers on a range of computer architectures,” *Journal of Computational Science*, Vol. 18, 2017, pp. 12 – 23. <https://doi.org/https://doi.org/10.1016/j.jocs.2016.11.001>.
- [17] Reguly, I. Z., Mudalige, G. R., Giles, M. B., Curran, D., and McIntosh-Smith, S., “The OPS Domain Specific Abstraction for Multi-block Structured Grid Computations,” *IEEE Press*, 2014, pp. 58–67. <https://doi.org/10.1109/WOLFHPC.2014.7>.
- [18] Shu, C.-W., “Essentially Non-Oscillatory and Weighted Essentially Non-Oscillatory Schemes for Hyperbolic Conservation Laws Operated by Universities Space Research Association,” *ICASE Report*, , No. 97-65, 1997, pp. 1–78.
- [19] Jiang, G.-S., and Shu, C.-W., “Efficient Implementation of Weighted ENO Schemes,” *Journal of Computational Physics*, Vol. 126, No. 1, 1996, pp. 202–228. <https://doi.org/https://doi.org/10.1006/jcph.1996.0130>.
- [20] Ducros, F., Ferrand, V., Nicoud, F., Weber, C., Darracq, D., Gacherieu, C., and Poinso, T., “Large-Eddy Simulation of the Shock/Turbulence Interaction,” *Journal of Computational Physics*, Vol. 152, No. 2, 1999, pp. 517 – 549. <https://doi.org/https://doi.org/10.1006/jcph.1999.6238>.
- [21] Ducros, F., Laporte, F., Soulères, T., Guinot, V., Moinat, P., and Caruelle, B., “High-Order Fluxes for Conservative Skew-Symmetric-like Schemes in Structured Meshes: Application to Compressible Flows,” *Journal of Computational Physics*, Vol. 161, No. 1, 2000, pp. 114 – 139. <https://doi.org/https://doi.org/10.1006/jcph.2000.6492>.

- [22] Sharma, N., and Sengupta, T. K., “Vorticity dynamics of the three-dimensional Taylor-Green vortex problem,” *Physics of Fluids*, Vol. 31, No. 3, 2019, p. 35106. <https://doi.org/https://doi.org/10.1063/1.5083870>.
- [23] Houba, T., Dasgupta, A., Gopalakrishnan, S., Gosse, R., and Roy, S., “Supersonic turbulent flow simulation using a scalable parallel modal discontinuous Galerkin numerical method,” *Scientific Reports*, Vol. 9, No. 1, 2019, p. 14442. <https://doi.org/https://doi.org/10.1038/s41598-019-50546-w>.
- [24] Sarkar, S., Erlebacher, G., Hussaini, M. Y., and Kreiss, H. O., “The analysis and modelling of dilatational terms in compressible turbulence,” *Journal of Fluid Mechanics*, Vol. 227, 1991, p. 473–493. <https://doi.org/https://doi.org/10.1017/S0022112091000204>.
- [25] Zeman, O., “Dilatation dissipation: The concept and application in modeling compressible mixing layers,” *Physics of Fluids A: Fluid Dynamics*, Vol. 2, No. 2, 1990, pp. 178–188. <https://doi.org/https://doi.org/10.1063/1.857767>.
- [26] Chong, M. S., Perry, A. E., and Cantwell, B. J., “A general classification of three-dimensional flow fields,” *Physics of Fluids A: Fluid Dynamics*, Vol. 2, No. 5, 1990, pp. 765–777. <https://doi.org/https://doi.org/10.1063/1.857730>.
- [27] Vreman, A. W., Sandham, N. D., and Luo, K. H., “Compressible mixing layer growth rate and turbulence characteristics,” *Journal of Fluid Mechanics*, Vol. 320, 1996, pp. 235–258. <https://doi.org/https://doi.org/10.1017/S0022112096007525>.
- [28] Sandham, N. D., “Effects of Compressibility and Shock-Wave Interactions on Turbulent Shear Flows,” *Flow, Turbulence and Combustion*, Vol. 97, No. 1, 2016, pp. 1–25. <https://doi.org/https://doi.org/10.1007/s10494-016-9733-6>.
- [29] Wang, Z., Fidkowski, K., Abgrall, R., Bassi, F., Caraeni, D., Cary, A., Deconinck, H., Hartmann, R., Hillewaert, K., Huynh, H., Kroll, N., May, G., Persson, P.-O., van Leer, B., and Visbal, M., “High-order CFD methods: current status and perspective,” *International Journal for Numerical Methods in Fluids*, Vol. 72, No. 8, 2013, pp. 811–845. <https://doi.org/10.1002/fld.3767>.
- [30] Mudalige, G. R., Reguly, I. Z., Giles, M. B., Mallinson, A. C., Gaudin, W. P., and Herdman, J. A., “Performance analysis of a high-level abstractions-based hydrocode on future computing systems,” *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, Vol. 8966, 2015, pp. 85–104. https://doi.org/10.1007/978-3-319-17248-4_5.
- [31] Mudalige, G. R., Reguly, I. Z., Jammy, S. P., Jacobs, C. T., Giles, M. B., and Sandham, N. D., “Large-scale performance of a DSL-based multi-block structured-mesh application for Direct Numerical Simulation,” *Journal of Parallel and Distributed Computing*, Vol. 131, 2019, pp. 130–146. <https://doi.org/10.1016/j.jpdc.2019.04.019>.
- [32] Lusher, D. J., Jammy, S. P., and Sandham, N. D., “OpenSBLI: Automated code-generation for heterogeneous computing architectures applied to compressible fluid dynamics on structured grids,” , 2020. arXiv:2007.14933.