

Positivity-preserving discontinuous spectral element methods for compressible multi-species flows

Will Trojak^{a,*}, Tarik Dzanic^b

^aIBM Research Europe, The Hartree Centre, Warrington, WA4 4AD, Cheshire, UK

^bDepartment of Mechanical and Aerospace Engineering, Princeton University, Princeton, 08544, NJ, USA

ARTICLE INFO

Keywords:

Compressible multi-species flows
High-order
Discontinuous spectral elements
Positivity-preserving
Entropy filtering

ABSTRACT

We introduce a novel positivity-preserving numerical stabilisation approach for high-order discontinuous spectral element approximations of compressible multi-species flows. The underlying stabilisation method is the adaptive entropy filtering approach (Dzanic and Witherden, *J. Comput. Phys.*, 468, 2022), which is extended to the conservative formulation of the multi-species flow equations. We show that the straightforward enforcement of entropy constraints in the filter yields poor results around species interfaces and propose an adaptive switch for the entropy bounds based on the convergence properties of the pressure field which drastically improves its performance for multi-species flows. The proposed approach is shown in a variety of numerical experiments applied to the multi-species Euler and Navier–Stokes equations computed on unstructured grids, ranging from shock-fluid interaction problems to three-dimensional viscous flow instabilities. We demonstrate that the approach can retain the high-order accuracy of the underlying numerical scheme even at smooth extrema, ensure the positivity of the species density and pressure in the vicinity of shocks and contact discontinuities, and accurately predict small-scale flow features with minimal numerical dissipation.


1. Introduction

The interaction and mixing of miscible fluids remain extensively studied phenomena, spanning various application areas such as combustion, climate modelling, and chemical process engineering [12, 30, 37, 42]. This behaviour is further complicated in specific applications with high-speed flows as fluid velocities are sufficiently large such that the compressibility effects cannot be ignored. The interaction of these multi-species flows with high Mach number effects, exhibiting features such as shocks, contact discontinuities, rarefaction waves, and species interfaces, poses a significant challenge for numerical modelling approaches. As such, the development of high-fidelity numerical methods for compressible multi-species flows remains an ongoing field of research with the potential for use in many practical applications.

Various models have been proposed for compressible multi-species flows, with the early work of Baer and Nunziato [6] presenting some competing approaches and introducing the now widely-adopted method of approximating the evolution of the density of each species along with the total momentum and energy. This conservative formulation can be straightforwardly implemented and is equally valid in both the incompressible, subsonic and compressible, supersonic regimes. While this approach suffers from the deficiency that conservative formulations cannot preserve pressure equilibrium across species interfaces [2, 3, 25, 26], its conservative nature ensures that it can accurately predict shock speeds and strong compressibility effects and, as a result, is critical for modelling multi-species flows in the high Mach regime. However, complex flow features in this class of flows present a challenge in developing efficient and accurate numerical approaches for these governing equations.

For obtaining these high-fidelity approximations of complex fluid flows, high-order discontinuous spectral element methods (DSEM) have grown in popularity, primarily due to their geometric flexibility, arbitrarily high-order accuracy, and compact data structure suited for high-performance computing. These benefits come at the expense of robustness issues, often in the form of numerical instabilities around discontinuous flow features commonly encountered in the approximation of compressible multi-species flows. The use of DSEM for this class of flows has been attempted in the works such as that of Billet and Ryan [8], Johnson and Kercher [26], Peyvan et al. [32], and Tonicello and

*Corresponding author

 w.trojak@ibm.com (W. Trojak)

ORCID(s): 0000-0002-4407-8956 (W. Trojak); 0000-0003-3791-1134 (T. Dzanic)

Ihme [41], and these approaches broadly rely on the use of additional numerical stabilisation techniques to ensure that the solution remains well-behaved in the vicinity of discontinuities. However, much like in the application of high-order DSEM to single-species gas dynamics and viscous fluid flows, it is notoriously difficult to design numerical stabilisation approaches that guarantee that the solution remains well-behaved while retaining the accuracy of the high-order numerical scheme, particularly so in a computationally efficient manner without the use of problem-dependent tunable parameters.

Therefore, this work aims to introduce and validate a robust and efficient numerical approach for simulating both inviscid and viscous compressible multi-species flows in the context of DSEM. This proposed approach is provably robust in that it guarantees that physical constraints such as positivity of density and pressure are satisfied while also retaining the high-order accuracy and scale-resolving capabilities of DSEM in smooth regions of the flow. Furthermore, it can be efficiently implemented on general unstructured meshes and does not require problem-dependent tunable parameters. The underlying stabilisation method of this approach is the entropy filter of Dzanic and Witherden [14], which enforces constraints such as positivity of density and pressure and a local minimum entropy principle through an adaptive filtering procedure. We show that straightforward implementation of the entropy constraints in the filter is exceedingly deficient for multi-species flows, particularly so around species boundaries, and yields poor results. We then present a novel method for adapting the entropy bounds based on the pressure field by utilising the convergence properties of DSEM through the approach of Krivodonova et al. [28] which significantly increases the accuracy of the stabilisation approach for multi-species flows. This proposed approach is validated and presented across a variety of multi-species problems on both structured and unstructured grids, ranging from simple transport and shock tube problems to inviscid shock-fluid interactions and three-dimensional viscous flow instabilities.

The remainder of this paper is organised as follows. In Section 2, we present some preliminaries on the governing equations and their properties as well as the underlying high-order DSEM. Then, we introduce in-depth the proposed numerical stabilisation approach in Section 3, followed by implementation details in Section 4. The results of the approach as applied to a variety of flow problems are then shown in Section 5, and conclusions are finally drawn in Section 6.

2. Preliminaries

2.1. Governing equations

Assuming thermal equilibrium, the governing equations are given as the evolution of n compressible species along with total momentum and energy. We use a species composition based on mass fraction similar to that of Abgrall [1]. In conservation form, the *inviscid limit* can be represented as

$$\frac{\partial}{\partial t} \begin{bmatrix} \boldsymbol{\alpha}\boldsymbol{\rho} \\ \rho\mathbf{V} \\ E \end{bmatrix} + \nabla \cdot \begin{bmatrix} \mathbf{V}^T \boldsymbol{\alpha}\boldsymbol{\rho} \\ \rho\mathbf{V} \otimes \mathbf{V} + \mathbf{I}P \\ \mathbf{V}(E + P) \end{bmatrix} = 0, \quad (1)$$

where

$$\boldsymbol{\alpha}\boldsymbol{\rho} = \begin{bmatrix} \alpha_0\rho_0 \\ \vdots \\ \alpha_{n-1}\rho_{n-1} \end{bmatrix} \quad \text{and} \quad \rho = \sum_{i=0}^{n-1} \alpha_i\rho_i \quad (2)$$

are the vector of species densities and the total density, respectively, $\rho\mathbf{V} \in \mathbb{R}^d$ is the total momentum, and E is the total energy. Furthermore, $\rho\mathbf{V}/\rho$ is the velocity and P is the pressure. To close this system, we utilise the equation of state

$$P = (\bar{\gamma} - 1)\rho e \quad \text{and} \quad \rho e = E - \frac{1}{2}\rho\mathbf{V} \cdot \mathbf{V}, \quad (3)$$

where ρe is the internal energy and $\bar{\gamma}$ is the specific heat ratio of the total mixture. Assuming thermal equilibrium, this specific heat ratio is calculated through the following mixture relation

$$\bar{\gamma} = \frac{\sum_{i=0}^{n-1} \alpha_i\rho_i c_{p,i}}{\sum_{i=0}^{n-1} \alpha_i\rho_i c_{v,i}}, \quad (4)$$

where $c_{p,i}$ and $c_{v,i}$ are the specific heat capacities at constant pressure and constant volume, respectively, of the given species. This contrasts with the approach of Ton [40] where the thermal equilibrium assumption is removed.

The inviscid governing equations can be readily extended to their viscous formulation as

$$\frac{\partial}{\partial t} \begin{bmatrix} \boldsymbol{\alpha}\boldsymbol{\rho} \\ \rho\mathbf{V} \\ E \end{bmatrix} + \nabla \cdot \begin{bmatrix} \mathbf{V}^T \boldsymbol{\alpha}\boldsymbol{\rho} \\ \rho\mathbf{V} \otimes \mathbf{V} + \mathbf{I}P \\ \mathbf{V}(E + P) \end{bmatrix} = \begin{bmatrix} 0 \\ \bar{\mu}(\nabla\mathbf{V} + \nabla\mathbf{V}^T) - \bar{\mu}\frac{2}{3}\nabla \cdot \mathbf{V} \\ \bar{\mu}(\nabla\mathbf{V} + \nabla\mathbf{V}^T)\mathbf{V} - \bar{\mu}\frac{\bar{\gamma}}{Pr}\nabla e \end{bmatrix}, \quad (5)$$

where $\bar{\mu}$ and \bar{Pr} are the dynamic viscosity and Prandtl number of the mixture, respectively, and $e = \rho e / \rho$ is the specific internal energy. There are several closure modules for the transport coefficients of the mixture [5, 20, 44]. In this work, we opt for a simple mass fraction weighting, i.e.,

$$\bar{\mu} = \frac{1}{\rho} \sum_{i=0}^{n-1} \alpha_i \rho_i \mu_i \quad \text{and} \quad \bar{Pr} = \frac{1}{\rho} \sum_{i=0}^{n-1} \alpha_i \rho_i Pr_i, \quad (6)$$

where Pr is the Prandtl number as this model gives good agreement with the method of Wilke [44].

2.2. Entropy principles

From Gouasmi et al. [19], for a specific choice of entropy functional σ , the inviscid (i.e., hyperbolic) governing equations satisfy an entropy inequality in the form of

$$\frac{\partial \rho\sigma}{\partial t} + \nabla \cdot (\mathbf{V}\rho\sigma) \geq 0, \quad (7)$$

where

$$\rho\sigma = \sum_{i=0}^{n-1} \alpha_i \rho_i \sigma_i \quad (8)$$

is the mass fraction weighted entropy. The entropy of each species is defined as

$$\sigma_i = \int_0^T \frac{c_{v,i}(\tau)}{\tau} d\tau - R_i \log(\rho_i), \quad (9)$$

where T is the temperature and R is the universal gas constant. Assuming that c_v does not vary with temperature, this mixture entropy can be expressed as

$$\sigma = \sum_{i=0}^{n-1} c_{v,i} \alpha_i \rho_i \log((\alpha_i \rho_i)^{1-\gamma_i} T), \quad (10)$$

where

$$T = \frac{\rho e}{\sum_{i=0}^{n-1} c_{v,i} \alpha_i \rho_i}. \quad (11)$$

Gouasmi et al. [19] showed that for any arbitrary \mathbf{x}_0 , this entropy (and by extent, all convex functionals of it) satisfies a local minimum entropy principle in the form of

$$\sigma(\mathbf{x}_0, t + \Delta t) \geq \min_{\mathbf{x} \in D_0} \sigma(\mathbf{x}, t) \quad (12)$$

for all $\Delta t \geq 0$, where D_0 is the domain of influence of \mathbf{x}_0 over the range $[t, t + \Delta t]$.

2.3. Discontinuous spectral element methods

The governing equations as described by Eq. (5) can be presented in the form of a general conservation law as

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

where \mathbf{u} is the vector of conserved variables and \mathbf{F} is the flux. We present here a very brief overview of the nodal discontinuous spectral element method (e.g., discontinuous Galerkin [22], flux reconstruction [24], etc.) as applied to first-order hyperbolic conservation laws, but for a more in-depth description and the extension to second-order systems, the reader is referred to Hesthaven and Warburton [22] and Huynh [24] (and the referenced works therein).

In this approach, the domain Ω is partitioned into N elements Ω_k such that $\Omega = \bigcup_{N_e} \Omega_k$ and $\Omega_i \cap \Omega_j = \emptyset$ for $i \neq j$. With a slight abuse of notation, the discrete solution $\mathbf{u}(\mathbf{x})$ within each element Ω_k can be formed through a nodal approximation as

$$\mathbf{u}(\mathbf{x}) = \sum_{i \in S} \mathbf{u}_i \phi_i(\mathbf{x}), \quad (14)$$

where $\mathbf{x}_i \forall i \in S$ is a set of solution nodes, $\phi_i(\mathbf{x})$ are their associated nodal basis functions, and S is the set of nodal indices for the stencil. The basis functions possess the property that $\phi_i(\mathbf{x}_j) = \delta_{ij}$. We utilise the notation that $\mathbf{u}_i = \mathbf{u}(\mathbf{x}_i)$ and that \mathbb{P}_p represents a p -th order approximation, taken as the maximal order of $\mathbf{u}(\mathbf{x})$. Furthermore, we assume that on the element interfaces $\partial\Omega_k$, there is a set of interface nodes $\mathbf{x}_i \in \partial\Omega \forall i \in \mathcal{I}$, where \mathcal{I} is a set of nodal indices for the interface stencil that is a subset of the solution nodes (i.e., $\mathcal{I} \subset S$). Due to the piecewise continuous nature of the solution approximation, there exist two values of the solution at each of these interface nodes, one from the element of interest, denoted as \mathbf{u}_i^- , and one from the interface-adjacent element, denoted as \mathbf{u}_i^+ .

A general formulation for the approximation of the flux $\mathbf{f}(\mathbf{x})$ in DSEM can be given as a collocation projection of the flux onto the solution nodes augmented with an interface correction term to account for inter-element interactions.

$$\mathbf{f}(\mathbf{x}) = \sum_{i \in S} \mathbf{F}(\mathbf{u}_i) \phi_i(\mathbf{x}) + \sum_{i \in \mathcal{I}} \bar{\mathbf{F}}(\mathbf{u}_i^-, \mathbf{u}_i^+, \mathbf{n}_i) \cdot \bar{\phi}_i(\mathbf{x}). \quad (15)$$

Here, $\bar{\mathbf{F}}(\mathbf{u}_i^-, \mathbf{u}_i^+, \mathbf{n}_i)$ denotes some common interface flux dependent on the interior/exterior values of the solution at the interfaces and their respective normal vector \mathbf{n}_i , which is commonly computed using exact or approximate Riemann solvers, e.g., Godunov [18], Rusanov [36], Roe [35]. Furthermore, $\bar{\phi}_i(\mathbf{x})$ denotes the correction basis function associated with the given interface node \mathbf{x}_i , which can be appropriately chosen to recover methods such as the discontinuous Galerkin approach (see Huynh [24]). With this approximation of the flux, the semi-discrete form of the governing equations, given as

$$\frac{\partial}{\partial t} \mathbf{u}_i = -\nabla \cdot \mathbf{f}(\mathbf{x}_i), \quad (16)$$

can be readily advanced in time using a suitable temporal integration algorithm.

3. Methodology

In this section, we introduce the proposed stabilisation method for high-order DSEM approximations of compressible multi-species flows. The underlying mechanisms of this approach rely on the properties of DSEM and the governing equations presented in Section 2. We assume here that the numerical discretisation is chosen such as to recover the nodal discontinuous Galerkin method [22] (either directly or via the flux reconstruction method [24]), strong-stability preserving (SSP) explicit time stepping is used, and the common interface fluxes are computed using a positivity-preserving entropy-stable Riemann solver. The method is first introduced with respect to inviscid compressible multi-species flows and extensions to their viscous counterparts are later presented.

3.1. Adaptive entropy filter

To stabilise the solution in the vicinity of discontinuities in a robust manner, Dzanic and Witherden [14] introduced a novel method of enforcing discrete constraints on the solution via an adaptive filtering procedure. We present here a

brief overview of the entropy filtering approach in the context of the original work in single-species gas dynamics and then show how this method can be extended to the multi-species system as well as what modifications and improvements are necessary. The goal of the approach is to ensure that

$$\mathbf{\Gamma}(\mathbf{u}_i) > 0 \quad (17)$$

is satisfied for all $i \in \mathcal{S}$, where $\mathbf{\Gamma}(\mathbf{u}) = \{\Gamma_1(\mathbf{u}), \Gamma_2(\mathbf{u}), \dots\}$ is some set of constraint functionals typically related to convex invariants of the system. To ensure robustness, physical constraints were enforced in the form of positivity of density and pressure (i.e., $\Gamma_1(\mathbf{u}) = \rho$, $\Gamma_2(\mathbf{u}) = P$). However, simply enforcing these physical constraints was not sufficient to ensure that the solution remains well-behaved in the vicinity of discontinuities. The key mechanism for stabilisation introduced by Dzanic and Witherden [14] was a forward-in-time local minimum entropy constraint of the form

$$\Gamma_3(\mathbf{u}) = \sigma - \sigma_{\min}, \quad (18)$$

where σ_{\min} is the discrete minimum entropy within the element and its Voronoi neighbours (see Dzanic and Witherden [14], Section 3.2).

These constraints were enforced through an adaptive modal filter, where the solution was first decomposed into its modal form as

$$\mathbf{u}(\mathbf{x}) = \sum_{i \in \mathcal{S}} \hat{\mathbf{u}}_i \psi_i(\mathbf{x}), \quad (19)$$

for some set of modal basis functions $\psi_i(\mathbf{x})$ computed with respect to the unit measure and their associated modes $\hat{\mathbf{u}}_i$. A filtered solution was then computed using a second-order exponential filtering kernel as

$$\tilde{\mathbf{u}}(\mathbf{x}) = \sum_{i \in \mathcal{S}} \hat{\mathbf{u}}_i e^{-\zeta p_i^2} \psi_i(\mathbf{x}), \quad (20)$$

where ζ is the filter strength and p_i is the maximal order of the basis function $\psi_i(\mathbf{x})$. This filtering operation can be considered as a form of hierarchical slope limiting as discussed in Michoski et al. [31]. The filter strength was computed adaptively through an element-wise scalar optimisation problem where the minimum necessary filter strength is sought such that the filtered solution satisfies the constraints discretely, i.e.,

$$\zeta = \arg \min_{\zeta \geq 0} \text{ s.t. } [\Gamma_1(\tilde{\mathbf{u}}(\mathbf{x}_i)) > 0, \Gamma_2(\tilde{\mathbf{u}}(\mathbf{x}_i)) > 0, \Gamma_3(\tilde{\mathbf{u}}(\mathbf{x}_i)) > 0 \quad \forall i \in \mathcal{S}]. \quad (21)$$

This optimisation problem can be readily solved using standard root-bracketing approaches, and due to the equivalency of the element-wise mean of discontinuous Galerkin schemes to first-order Godunov methods, a solution to this optimisation problem is guaranteed to exist given the previously mentioned assumptions on the numerical scheme [10, 14, 46–48].

The extension of this approach to the multi-species system presents additional options for constraints on the solution. First, the physical constraints on the positivity of the (mixture) density and pressure are retained, i.e.,

$$\Gamma_1(\mathbf{u}) = \rho \quad \text{and} \quad \Gamma_2(\mathbf{u}) = P. \quad (22)$$

Then, further constraints are introduced for the positivity of the individual species densities as

$$\Gamma_3(\mathbf{u}) = \alpha_i \rho_i \quad (23)$$

for $0 \leq i \leq n-1$. Note that while from an analytical perspective, $\Gamma_1(\mathbf{u})$ is redundant given $\Gamma_3(\mathbf{u})$, it is advantageous to separate these constraints numerically due to the inclusion of numerical tolerances. This separation is later discussed in Section 4. Finally, a straightforward application of the entropy filter of Dzanic and Witherden [14] can be done by including the entropy constraint as

$$\Gamma_4(\mathbf{u}) = \sigma - \sigma_{\min}, \quad (24)$$

where σ is defined by Eq. (10).

However, as will be shown in the numerical experiments, the naive implementation of the entropy constraints for compressible multi-species flows results in exceedingly deficient approximations, particularly around species interfaces. This deficiency stems from the discontinuities introduced by species interfaces, where even flows with smooth mixture density fields may consist of individually discontinuous species density fields which, by Eq. (9) and Eq. (10), can cause numerically ill-behaved entropy fields. As a result, excessive numerical dissipation is introduced in the vicinity of species interfaces which, much like contact discontinuities in single-species flows, do not need nearly as much stabilisation as compressive features such as shocks.

A novelty of the proposed approach is the introduction of a method to modify these entropy constraints such that they can adequately stabilise the solution in regions where it is necessary while retaining high accuracy around species interfaces and contact discontinuities. To retain the benefits of the entropy filtering approach, this modification must maintain its desirable properties, namely the efficiency, robustness, and lack of problem-dependent tunable parameters. To this end, we propose a switch to the entropy bounds inspired by the work of Gao et al. [17], such that the constraint can be expressed as

$$\Gamma_4(\mathbf{u}) = \begin{cases} \sigma - \sigma_{\min}, & \text{if } S \geq 1, \\ 0 & \text{else,} \end{cases} \quad (25)$$

where S is some element-wise switching function. The goal of this switching function is to be able to isolate regions where the numerical scheme would exhibit numerical instabilities (i.e., shocks) such that the entropy constraints are necessary to ensure a well-behaved solution while neglecting other discontinuous regions such as contact discontinuities and species interfaces. Most importantly, this switching function should be able to achieve this in a way such as to not require problem-dependent tunable parameters. We note here that regardless of the choice of S , *the physical constraints are always enforced*, such that the positivity-preserving properties are independent of the switching function and its purpose is purely to increase the accuracy of the numerical scheme around species interfaces.

3.2. Interface jump convergence sensor

The proposed approach relies on detecting discontinuities in the pressure field to isolate contact discontinuities and species interfaces from shocks in the flow. While many methods exist for sensing discontinuities in the solution, they typically rely on arbitrary tunable parameters or are difficult to extend to high-order DSEM on unstructured grids in a robust and efficient manner. To this end, we choose to formulate the switching function similarly to the method of Krivodonova et al. [28] which utilises the convergence properties of discontinuous Galerkin-type methods to isolate discontinuities in the solution within elements. For DSEM, it is known that the magnitude of the jump in the solution (or some derived quantity thereof) across an element interface is dependent on the regularity of the solution. If we define the jump of some derived quantity of the solution Q across $\partial\Omega_i$ as

$$I_i = \int_{\partial\Omega_i} (Q_i^+ - Q_i^-) ds, \quad (26)$$

then it is expected that I_i is of at least $\mathcal{O}(h^{p+1})$, where h is some characteristic mesh length, if the solution is sufficiently smooth in the vicinity of Ω_i . However, if the solution is discontinuous in the vicinity of Ω_i , then I_i would typically be of $\mathcal{O}(h)$ (or higher if the discontinuity is aligned with the element interface). By utilising the convergence properties of the pressure field, the switching function can be defined as

$$S_i = \frac{\left| \int_{\partial\Omega_i} (P_j^+ - P_j^-) ds \right|}{h^{p+1} \|P_j^-\|_{\infty} |\partial\Omega_i|}. \quad (27)$$

With this formulation, the switching function can be evaluated and the entropy bounds can be appropriately modified per Eq. (25) at the time of the filter application. We remark here that the threshold value of unity is not necessarily a strict condition as some arbitrariness in this value exists. However, this value was chosen for consistency with the original work of Krivodonova et al. [28], and, as will be shown, this choice yields a robust and accurate stabilisation approach for a series of challenging numerical experiments. Furthermore, as previously mentioned, the positivity-preserving properties of the scheme are independent of the choice of threshold value. The numerical procedure for evaluating the switching function as well as the formulation for computing the characteristic mesh scale for general grids is presented in Section 4.

3.3. Operator splitting

As the proposed scheme is reliant on the well-posedness of a local minimum entropy principle, care must be taken in its extension to the multi-species Navier–Stokes equations. While the single-species counterpart of the viscous flow equations may satisfy a minimum entropy principle for a particular choice of entropy [39], it is not evident whether this holds for the multi-species case with the given entropy functional nor whether the numerical scheme would satisfy the minimum entropy principle on the element-wise mean in the presence of the viscous dissipation. To use the proposed scheme for the Navier–Stokes equations, two options can be considered. One may simply enforce entropy constraints on the full hyperbolic–parabolic system even with the possibility of there not existing a solution to the element-wise optimisation problem. Note that this option does not negatively impact the positivity-preserving properties and robustness of the scheme as any stable root-bracketing method would converge to the maximal filter strength in the absence of a solution to the optimisation problem.

The alternative approach, which is considered in this work, is to split the hyperbolic (inviscid) and parabolic (viscous) operators and selectively apply constraints as necessary. The semi-discrete form in Eq. (16) is represented in terms of the inviscid and viscous flux as

$$\partial_t \mathbf{u} + \nabla \cdot (\mathbf{F}_I(\mathbf{u}) + \mathbf{F}_V(\mathbf{u})) = 0. \quad (28)$$

For some arbitrary n and time step $\Delta t > 0$, an intermediate temporal update can be given as

$$\bar{\mathbf{u}}^{n+1} = \mathbf{u}^n - \Delta t \nabla \cdot (\mathbf{F}_I(\mathbf{u}^n)), \quad (29)$$

corresponding to an inviscid substep of the governing equations. If necessary, the boundary conditions must be appropriately modified to ensure consistency with the governing equations (e.g., no slip wall boundary conditions must be converted to slip wall boundary conditions). As this substep obeys the local discrete minimum entropy principle, the adaptive filtering operation $H(\mathbf{u})$ can be applied as usual to stabilise the solution where necessary. Afterwards, the full temporal update can be performed by adding the viscous component of the divergence of the flux as

$$\mathbf{u}^{n+1} = H(\bar{\mathbf{u}}^{n+1}) - \Delta t \nabla \cdot (\mathbf{F}_V(\mathbf{u}^n)). \quad (30)$$

Note that the viscous flux is evaluated in an explicit manner, not through a Strang splitting-type approach, such that the accuracy of the temporal integration method is not detrimentally affected. While the addition of the viscous component typically further stabilises the solution, it is possible in rare scenarios that positivity of density and pressure may not be ensured for \mathbf{u}^{n+1} even though it is for $H(\bar{\mathbf{u}}^{n+1})$. Therefore, a secondary filtering operation is applied to the \mathbf{u}^{n+1} using only positivity constraints to guarantee that the solution remains positivity-preserving. As the violation of the positivity constraints caused by the additional viscous components is very rare [14], in the vast majority of situations, the secondary filtering operation is not performed.

The operator splitting approach is presented in terms of a single step of a forward Euler scheme. The approach can naturally be extended to any explicit SSP temporal integration method in a similar manner. For extensions to higher-order SSP schemes as well as a detailed overview of the operator splitting approach, the reader is referred to Dzanic and Witherden [14], Appendix A and Dzanic and Witherden [15], Section 3.2.

4. Implementation

The proposed numerical scheme was implemented within PyFR [45], a high-order flux reconstruction solver that can be efficiently deployed on massively-parallel GPU and CPU computing architectures. For a given polynomial approximation order p , the solution and flux points were distributed along the corresponding Gauss–Legendre–Lobatto quadrature points for tensor-product elements and α -optimised [22] points for simplex elements. Temporal integration was performed using a third-order, three stage SSP Runge–Kutta scheme with a fixed time step. The common inviscid interface fluxes were computed using the HLLC Riemann solver [4, 43] as it was found to yield notably better results than approaches such as that of Rusanov [36]. For the common viscous interface fluxes, the BR2 approach of Bassi and Rebay [7] was used.

To ensure a non-vacuum state for the Riemann solver, a numerical tolerance ϵ was added to the mixture density and pressure constraints, i.e.,

$$\Gamma_1(\mathbf{u}) = \rho - \epsilon, \quad \Gamma_2(\mathbf{u}) = P - \epsilon. \quad (31)$$

Furthermore, the same tolerance was applied to the entropy constraints as

$$\Gamma_4(\mathbf{u}) = \sigma - \sigma_{\min} + \epsilon. \quad (32)$$

Unless otherwise stated, this tolerance was set as $\epsilon = 10^{-5}$ in this work. For the constraint on the individual species density, the tolerance was not applied through the constraints themselves but instead through adding ϵ to the initial conditions for each species density component.

Due to the logarithmic nature of Eq. (10), the calculation of the entropy is highly prone to numerical precision issues as well as numerically undefined behaviour in the limit of zero species density. The robustness and accuracy of the proposed approach was significantly improved through the use of the modified entropy functional σ^* , defined as

$$\sigma^* = \exp \left[\sum_{i=0}^{n-1} c_{v,i} \alpha_i \rho_i \log \left(\max(\epsilon, \alpha_i \rho_i)^{1-\gamma_i T} \right) \right], \quad (33)$$

which was used for computing the entropy constraints in place of Eq. (10). As the exponential function is strictly convex, this modified entropy retains the local minimum entropy principle of its original form, and the clipping of the species density ensures that the entropy is numerically well-defined.

Unless otherwise stated, the numerical experiments are assumed to be computed with the entropy switch enabled. The sensor for the entropy switch, defined by Eq. (27), was computed discretely across the nodal interface points for each element Ω_i as

$$S_i = \frac{\sum_{j \in \mathcal{I}} w_j (P_j^+ - P_j^-)}{h^{p+1} \left(\max_{j \in \mathcal{I}} |P_j^-| \right) \left(\sum_{j \in \mathcal{I}} w_j \right)}, \quad (34)$$

where w_j is the corresponding quadrature weight for the given interface point \mathbf{x}_j on $\partial\Omega_i$. The formulation of the mesh scale h was chosen in such a way to easily generalise between unstructured and structured meshes of varying dimensionality. We define the mesh scale as the diameter of the d -ball of equivalent surface area to the element in question, i.e.,

$$h = \begin{cases} A/\pi, & \text{if } d = 2, \\ \sqrt{A/\pi}, & \text{if } d = 3, \end{cases} \quad (35)$$

where $A = \sum_{j \in \mathcal{I}} w_j$ is the surface area of $\partial\Omega_i$. While other approaches exist for computing a mesh scale for an unstructured mesh (e.g., inscribed circle/sphere), it was found that the results were not particularly sensitive to how the mesh scale was computed.

For the adaptive filtering method, the constraints were first checked each time the filtering operation was called, and if the solution satisfied the constraints, no filtering was applied. In the case that the solution violated the constraints, the filter strength was computed using the Illinois root-bracketing approach [11] with a maximum of 20 iterations and an early stopping tolerance of 10^{-8} . A highly-efficient optimisation procedure was utilised which exploits the structured nature of the filtering matrices (see Dzanic and Witherden [15], Section 4.1), such that the overall cost of the adaptive filtering method was only a small portion of the total compute time. This approach was implemented and deployed on parallel GPU computing architectures with computations performed on up to 40 NVIDIA V100 GPUs. Further details of the configurations used throughout this work can be found in the electronic supplementary material.

5. Results

5.1. Multi-species Euler equations

5.1.1. Near-vacuum convecting density wave

As an initial verification of the ability of the proposed approach to recover the high-order accuracy of the underlying DSEM for smooth problems, the convergence of the scheme was evaluated for a smooth convecting density wave. This problem is evaluated on the periodic domain $\Omega = [-0.5, 0.5] \times [-0.05, 0.05]$, and the initial conditions are given as

$$\rho = \alpha_0 \rho_0 + \alpha_1 \rho_1 = \exp(-\sigma x^2) + 4\epsilon, \quad (36a)$$

Table 1

Convergence in the L^∞ norm of the density error at $t = 1$ with respect to mesh resolution N for the near vacuum convection problem with varying approximation order and *with* the entropy switch. Rate of convergence shown beneath.

N	\mathbb{P}_1	\mathbb{P}_2	\mathbb{P}_3	\mathbb{P}_4	\mathbb{P}_5
8	-	-	5.32×10^{-1}	5.20×10^{-1}	5.31×10^{-1}
16	-	3.07×10^{-1}	6.94×10^{-2}	4.54×10^{-2}	2.14×10^{-1}
32	3.16×10^{-1}	3.36×10^{-2}	1.1×10^{-2}	3.82×10^{-4}	7.70×10^{-4}
64	1.04×10^{-1}	2.85×10^{-3}	1.91×10^{-4}	1.12×10^{-5}	8.05×10^{-7}
128	2.21×10^{-2}	3.53×10^{-4}	1.33×10^{-5}	5.41×10^{-7}	-
256	3.68×10^{-3}	4.35×10^{-5}	1.06×10^{-6}	-	-
512	6.26×10^{-4}	5.47×10^{-6}	-	-	-
RoC	2.28	3.16	3.93	5.17	6.61

$$\alpha_0 = 1 - \alpha_1 = \frac{1}{2}(\sin(2\pi x) + 1), \quad (36b)$$

$$u = 1, \quad (36c)$$

$$v = 0, \quad (36d)$$

$$P = 2\epsilon, \quad (36e)$$

where $\sigma = 500$ is the strength of the Gaussian wave. We consider the case of two-species, where the specific heat capacities were set as $c_{p,0} = 1.4$, $c_{v,0} = 1$, $c_{p,1} = 4.21$, and $c_{v,1} = 2.52$. To make this problem more challenging and to test the robustness of the positivity-preserving properties of the proposed scheme, a near-vacuum state was used for the initial conditions by setting $\epsilon = 10^{-12}$, such that each species density had a minimum value of $2 \cdot 10^{-12}$. The tolerances and minimum density/pressure conditions for the scheme were also modified accordingly. For these extreme conditions, numerical undershoots can easily cause negative density and pressure states which can cause a solver to diverge if it does not explicitly enforce positivity.

To evaluate the behaviour of the scheme at smooth extrema, the L^∞ norm of the density error was calculated after one flow-through of the domain. The error and the average rate of convergence for varying approximation orders and mesh resolution is shown in Table 1. It can be seen that the high-order accuracy of DSEM was not detrimentally affected by the proposed stabilisation method, such that high-order (i.e., $p + 1$) convergence was observed even at smooth extrema. However, it may be argued that the observed high-order accuracy for the given problem could be a by-product of the entropy switch as the constant pressure field would not trigger the entropy constraints (i.e., $S \sim 0$) which are primarily responsible for stabilisation in the vicinity of discontinuities. To verify that the stabilisation technique recovers high-order accuracy even with entropy constraints, the convergence tests were repeated with the entropy switch disabled, such that the entropy constraints were always enforced for every element in the domain. The error and convergence rates with the entropy switch disabled are presented in Table 2. It can clearly be seen that the resulting errors and, by extension, convergence rates, are nearly identical with and without the entropy switch. These results indicate that the proposed stabilisation approach retains the high-order accuracy of the underlying DSEM in smooth regions of the flow while ensuring that physical constraints such as positivity of density and pressure are satisfied by the discrete solution.

5.1.2. Isentropic Euler vortex

The convecting Euler vortex [38] is an isentropic solution to the Euler equations with an analytic expression for its time evolution. Consequently, it is frequently used as a validation case for numerical solvers. For the extension to multi-species flows, the initial conditions can be adapted as

$$\rho_i = \left(1 - \frac{(\gamma_i - 1)\beta^2 M^2}{8\pi^2} \exp(2f) \right)^{1/(\gamma_i - 1)}, \quad \text{for } i \in \{1, 2\}, \quad (37a)$$

$$u = U_0 + \frac{\beta y}{2\pi R} \exp(f), \quad (37b)$$

$$v = V_0 - \frac{\beta x}{2\pi R} \exp(f), \quad (37c)$$

Table 2

Convergence in the L^∞ norm of the density error at $t = 1$ with respect to mesh resolution N for the near vacuum convection problem with varying approximation order and *without* the entropy switch. Rate of convergence shown beneath.

N	\mathbb{P}_1	\mathbb{P}_2	\mathbb{P}_3	\mathbb{P}_4	\mathbb{P}_5
8	-	-	5.32×10^{-1}	5.20×10^{-1}	5.31×10^{-1}
16	-	3.07×10^{-1}	6.94×10^{-2}	4.54×10^{-2}	2.14×10^{-1}
32	3.16×10^{-1}	3.37×10^{-2}	1.1×10^{-2}	3.83×10^{-4}	7.70×10^{-4}
64	1.04×10^{-1}	2.85×10^{-3}	1.91×10^{-4}	1.11×10^{-5}	1.11×10^{-6}
128	2.21×10^{-2}	3.53×10^{-4}	1.32×10^{-5}	4.62×10^{-7}	-
256	3.68×10^{-3}	4.35×10^{-5}	1.1×10^{-6}	-	-
512	6.26×10^{-4}	5.46×10^{-6}	-	-	-
RoC	2.28	3.16	3.93	5.22	6.47

$$P = \frac{1}{\gamma M^2} \rho^\gamma, \quad (37d)$$

$$\alpha_1 = \frac{1}{2} \sin\left(\frac{\pi}{L}(y - y_0)\right) + \frac{1}{2}, \quad (37e)$$

$$f = \frac{1 - (x - x_0)^2 - (y - y_0)^2}{2R^2}, \quad (37f)$$

defined on the domain $\Omega = [x_0 - L, x_0 + L] \times [y_0 - L, y_0 + L]$. Here, the parameters $\beta = 13.5$ denote the strength of the vortex, $R = 1.5$ the radius, $U_0 = 0$, $V_0 = 1$ the advection velocities, and $M = 0.4$ the free-stream Mach number. The domain was centred around $x_0 = y_0 = 0$ with an extent of $L = 10$. In the multi-species case, forming an exact solution at an arbitrary time is complex. Therefore, we instead use a different procedure to inspect the efficacy of the proposed approach. First, the solution is advanced to some time $t = T$, after which the flow velocities are reversed and then advanced to $t = 2T$. If the initial condition is isentropic, then the solution should be reversible and any irreversibilities stemming from numerical error and dissipation would appear as discrepancies between the solutions at $t = 0$ and $t = 2T$. To yield this isentropic initial condition, we consider the two-species case with $\gamma_0 = \gamma_1$, recovered through the specific heat capacities $c_{p,0} = 1.4$, $c_{v,0} = 1$, $c_{p,1} = 2.8$, and $c_{v,1} = 2$.

The contours of the first species density at $t = T = 20$, computed by a \mathbb{P}_4 approximation with varying mesh resolution, are shown in Fig. 1 with and without the entropy switch. It can evidently be seen that there is a marked difference in the accuracy of the method when the entropy switch is used that is most pronounced in the vicinity of species interfaces. With the switch enabled, very good resolution of the vortex rollup was observed even at $N = 40^2$, showing, at least in the qualitative sense, mesh convergence with respect to the more resolved case of $N = 80^2$. However, when the entropy switch was disabled, the enforcement of entropy constraints around species interfaces had a severely detrimental effect on the accuracy of the proposed method, with poor convergence with respect to mesh resolution.

After flow reversal at $t = 2T$, shown in Fig. 2, the difference between the accuracy of the method with and without the entropy switch became even more pronounced. While the approach with the entropy switch was able to recover the initial conditions reasonably well with $N = 40^2$ and very well with $N = 80^2$, the approach without the entropy switch showed noticeable numerical artefacts at all resolutions as a result of the excessive numerical dissipation introduced by the entropy constraints at species interfaces. As a result, the predictions at $N = 80^2$ without the entropy switch were on the order of the much less resolved $N = 20^2$ simulation that did utilise the entropy switch. These observations are highlighted in Fig. 3 which shows a cross-section of the first species density at $x = 0$ at both the flow reversal time and final time computed with the entropy switch. It can be seen that the approach can resolve the highly-oscillatory nature of the vortex rollup and recover the initial conditions very well. Furthermore, a comparison of the results with and without the entropy switch is shown on the same cross-section in Fig. 4. The substandard resolving ability for the vortex rollup and the flow irreversibilities introduced by the significant numerical dissipation of the approach without the entropy switch can be clearly observed. These results showcase both the resolving ability of the scheme and the importance of adapting the entropy bounds by the proposed approach.




Figure 1: Contours of $\alpha_0 \rho_0$ for the multi-species isentropic Euler vortex problem at $t = T = 20$ using a \mathbb{P}_4 approximation with $N = 20^2$ (left), 40^2 (middle), and 80^2 (right) elements. Results with and without the entropy switch shown on top and bottom rows, respectively.

Table 3

Initial conditions for the Quirk and Karni shock tube problem.

Region	Extent	ρ_0	ρ_1	u	P	γ	c_v
L	$x \leq 0.25$	1.3765	0	0.3948	1.57	1.4	0.72
R	$0.25 < x \leq 0.4, x > 0.6$	1	0	0	1	1.4	0.72
B	$0.4 < x \leq 0.6$	0	0.138	0	1	1.67	2.52

5.1.3. Quirk and Karni shock tube

To evaluate the proposed scheme for more complex flow physics including shock waves, contact discontinuities, and rarefaction waves, the one-dimensional shock tube problem of Quirk and Karni [33], a simplification of a shock-bubble interaction, was considered. In this test case, the various features of the Riemann problem are observed as well as discontinuities between species, which can introduce pressure oscillations in conservative approaches where pressure equilibrium is not guaranteed. The problem is solved in a quasi-one-dimensional form on the domain $\Omega = [0, 1] \times [-0.05, 0.05]$, and the initial conditions are given by Table 3.

For this experiment, a \mathbb{P}_3 approximation was used with varying mesh resolution. The resulting density and pressure fields at $t = 0.35$ are shown in Fig. 5. A reference solution was computed with a highly resolved, first-order Godunov [18] scheme with $N = 5 \cdot 10^4$ elements. It can be seen that the proposed approach shows good convergence to the reference solution with increasing mesh resolution, with excellent agreement at $N = 200$. Furthermore, better

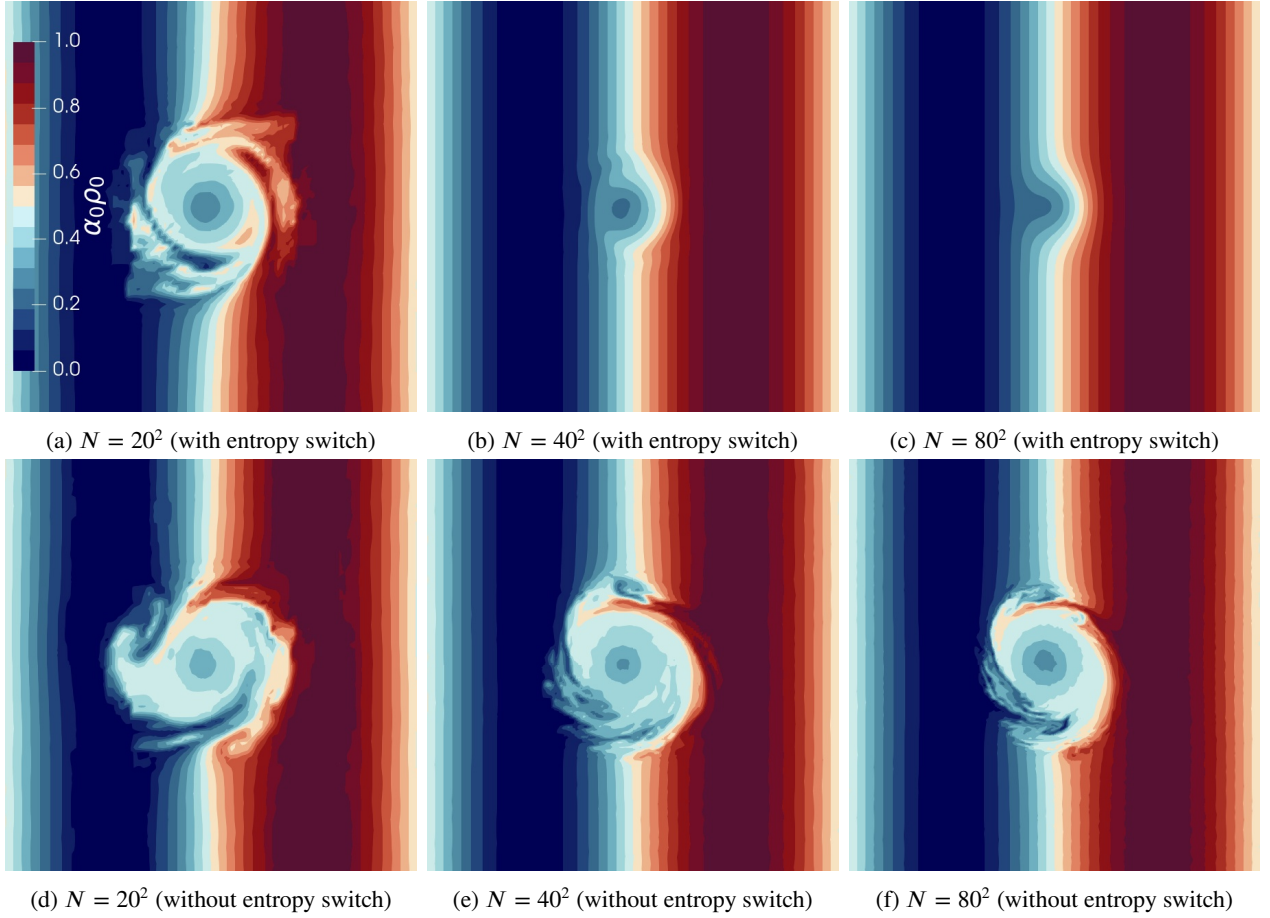


Figure 2: Contours of $\alpha_0\rho_0$ for the multi-species isentropic Euler vortex problem after flow reversal at $t = 2T = 40$ using a \mathbb{P}_4 approximation with $N = 20^2$ (left), 40^2 (middle), and 80^2 (right) elements. Results with and without the entropy switch shown on top and bottom rows, respectively.




Figure 3: Cross-section of $\alpha_0\rho_0$ at $x = 0$ for the multi-species isentropic Euler vortex problem at $t = T = 20$ (left) and after flow reversal at $t = 2T = 40$ (right) using a \mathbb{P}_4 approximation with varying mesh resolution computed *with* the entropy switch. Initial $\alpha_0\rho_0$ field at $t = 0$ shown by grey markers.




Figure 4: Comparison of the cross-section of $\alpha_0 \rho_0$ at $x = 0$ for the multi-species isentropic Euler vortex problem at $t = T = 20$ (left) and after flow reversal at $t = 2T = 40$ (right) using a \mathbb{P}_4 approximation with $N = 80^2$ computed with the entropy switch (red) and without the entropy switch (black). Initial $\alpha_0 \rho_0$ field at $t = 0$ shown by grey markers.

prediction of discontinuities in the flow was observed with increasing resolution, with less numerical dissipation introduced by scheme. In the pressure field, the oscillations associated with the pressure equilibrium problem [2], commonly seen in conservative approaches in the region $0.4 < x < 0.7$, were not visually noticeable. It is expected that these oscillations reduce with increasing mesh resolution [2, 26], such that their effects are minimal for well-resolved simulations.




Figure 5: Profiles of density (left) and pressure (right) for the Quirk and Karni shock tube problem at $t = 0.35$ computed using a \mathbb{P}_3 approximation with varying mesh resolution.

5.1.4. Helium-air shock-bubble interaction

The extension to more complex flow physics and unstructured meshes was performed through the helium-air shock-bubble interaction problem, a well-studied case first experimentally investigated by Haas and Sturtevant [21] and numerically simulated by Quirk and Karni [33]. The interaction between the travelling shock wave and the helium gas bubble seeds Richtmyer–Meshkov instabilities at the interface, yielding complex vortical flow. Accurately predicting this interaction is the primary challenge in this test case as stabilisation methods with excessive numerical dissipation tend to smear the interface and associated small-scale flow features.




Figure 6: Schematic of the domain and initial conditions for the helium-air shock-bubble interaction problem.

Table 4

Initial conditions for the helium-air shock-bubble interaction problem.

Region	ρ_0	ρ_1	u	v	P	γ	c_v
L	1	0	0	0	1	1.3986	0.72
R	1.37636	0	-0.55957	0	1.5698	1.3986	0.72
B	0	0.18187	0	0	1	1.6467	2.44

The numerical setup consists of a bubble of radius $R = 5/89$ filled with helium contaminated with 28% air by mass. This bubble is centred at $x = 0.8$, $y = 0$ within a channel of unit length and height of 0.2 filled with air. The resulting numerical domain is set as $\Omega = [0, 1] \times [-0.1, 0.1]$, shown in Fig. 6. At $x = 0.9$, a normal shock corresponding to a Mach number of $M = 1.22$ is imposed. The flow is normalised to unit density and pressure by the post-shock air state, denoted by the subscript l . The pre-shock air state, denoted by the subscript r , is computed via the Rankine–Hugoniot conditions. The bubble state, denoted with the subscript b , is set to be in thermal and mechanical equilibrium with the surrounding air. The initial conditions as well as the necessary gas constants for the three states are presented in Table 4.




Figure 7: Numerical Schlieren diagrams for the helium-air shock-bubble interaction at $t = 0.25$ computed on a coarse (left), medium (middle), and fine (right) unstructured triangular mesh with a \mathbb{P}_3 approximation. Solution on the half-domain is reflected across the x -axis for visualisation.

The problem was solved using a \mathbb{P}_3 approximation on unstructured meshes. Due to the symmetry of the problem, only half of the domain ($y > 0$) was solved for, with symmetry (slip adiabatic wall) boundary conditions imposed along the lower wall. Similarly, slip adiabatic wall boundary conditions were used for the upper wall. Dirichlet boundary conditions were used for the inlet ($x = 1$) and Neumann boundary conditions for the outlet ($x = 0$). A series of unstructured triangular meshes were generated, with the coarse, medium, and fine meshes corresponding to average edge lengths of $h/R = 50$, 100, and 200, respectively. The resulting numerical Schlieren diagrams at $t = 0.25$, showing the magnitude of the density gradient, are shown in Fig. 7 for the various mesh resolutions. The rollup of vortices along

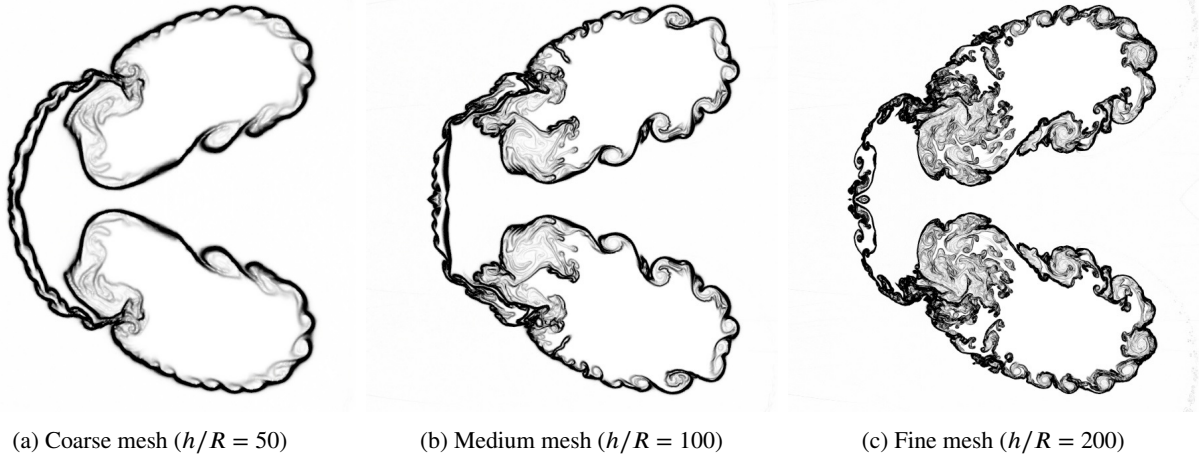


Figure 8: Numerical Schlieren diagrams for the helium-air shock-bubble interaction at $t = 0.39$ computed on a coarse (left), medium (middle), and fine (right) unstructured triangular mesh with a \mathbb{P}_3 approximation. Solution on the half-domain is reflected across the x -axis for visualisation.

the bubble interface were clearly observed, with smaller scale flow features appearing with increasing resolution. At a later time of $t = 0.39$, shown in Fig. 8, the bubble showed more lateral spread and distinct vortical structures, and similar observations with respect to the mesh resolution were drawn.




Figure 9: Comparison of the numerical Schlieren diagrams for the helium-air shock-bubble interaction between the present work (top) and Kundu [29] (bottom) at $t = 0.25$ (left) and $t = 0.39$ (right). Corresponding simulation time of Kundu [29] shown in parentheses.

A comparison of the proposed approach on the fine mesh and the results of Kundu [29], computed with a ninth-order upwind finite difference scheme, is shown in Fig. 9 at the two simulation times. Between the two methods, the effective mesh resolution in terms of solution point spacing was roughly similar. Very good agreement between the two approaches was observed, with near identical predictions of the dominant flow structures in the problem as well as good agreement in the shape and size of the small-scale flow features.

5.1.5. Triple-point shock interaction

As a final inviscid test case, the triple-point shock interaction problem of Galera et al. [16] was used to evaluate the method for complex shock-vortex interactions. The problem consists of a two-dimensional, two-species Riemann problem, with an overpressure state driving the rollup of a contact discontinuity between species. The domain is set as

$\Omega = [0, 7] \times [0, 3]$ with slip adiabatic walls on all four sides, and the initial conditions for the three states are shown in Fig. 10. The initial velocity is set to zero throughout the domain. To achieve the desired specific heat ratios of $\gamma_0 = \gamma_t = \gamma_b = 1.5$ and $\gamma_1 = \gamma_b = 1.4$, we simply set $c_{v,0} = c_{v,1} = 1$ and $c_p = \gamma$, respectively.




Figure 10: Schematic of the domain and initial conditions for the triple-point shock interaction problem.

The problem was solved using a \mathbb{P}_3 approximation on a uniform quadrilateral mesh with $N = 2800 \times 1200$ elements. The contours of mixture density at evenly-spaced time intervals are shown in Fig. 11. It can be seen that the rollup of the contact discontinuity introduces Kelvin–Helmholtz instabilities in the flow which were well-resolved by the proposed scheme. The interaction of these vortical structures with the reflecting shock waves in the domain showcases the ability of the scheme to stabilise the solution around discontinuities without excessively dissipating small-scale flow structures. These results compare favourably to works such as that of Kolev and Rieben [27] in terms of the resolving capability of small-scale flow structures.

5.2. Multi-species Navier–Stokes equations

5.2.1. Rayleigh–Taylor instability

The extension to the multi-species viscous flow equations was performed through simulations of the Rayleigh–Taylor instability, a fundamental phenomenon stemming from the interaction of fluids of different densities in hydrostatic equilibrium. The problem consists of a denser fluid resting above a lighter fluid, with gravitational forces initially balanced by an equivalent pressure gradient. The interface between the fluids forms instabilities manifesting as “fingers” of heavy fluid descending in conjunction with “pillars” of lighter fluid ascending, driving the flow into a chaotic mixing state.

The problem is solved on the domain $\Omega = [-0.5, 0.5]^2 \times [-1, 1]$. The fluid interface was placed at $z = 0$, with $\alpha_0 \rho_0 = \rho_b$, $\alpha_1 \rho_1 = 0$ for $z \leq 0$ and $\alpha_0 \rho_0 = 0$, $\alpha_1 \rho_1 = \rho_t$ for $z > 0$. The densities of the light and heavy fluid were taken as $\rho_b = 1$ and $\rho_t = 3$, respectively, yielding an Atwood number of $1/2$. The fluid was taken to be at rest in the transverse directions, i.e., $u = v = 0$, and a small vertical velocity perturbation was imposed to seed instabilities in the flow. A deterministic single-mode sinusoidal perturbation was used with a cosinusoidal vertical decay as

$$w = A \cos\left(\frac{\pi z}{2L}\right) \sin\left(\frac{4\pi x}{L}\right) \sin\left(\frac{4\pi y}{L}\right), \quad (38)$$

where the amplitude of the perturbation was taken as $A = 0.05$. Constant source terms of $-\rho g$ and $-\rho g w$, where $g = 1$, in the vertical momentum and energy equations, respectively, were used to simulate a gravitational field effect. The initial pressure field was chosen to enforce hydrostatic equilibrium in the flow with respect to the gravitational field, i.e.,

$$P = P_0 - \rho g z. \quad (39)$$

The ambient pressure was set as $P_0 = 6$.

A uniform hexahedral mesh with $N = 32 \times 32 \times 64$ elements was used with periodicity enforced along the transverse directions, and a slip adiabatic wall boundary condition was applied to the top and bottom walls. The specific heat constants for the lighter and heavier fluids were taken as $c_{v,0} = 3.11$, $c_{p,0} = 5.19$, $c_{v,1} = 0.72$, and $c_{p,1} = 1.007$. A uniform Prandtl number and dynamic viscosity were used, with the Prandtl number fixed at $Pr = 0.71$ and the dynamic viscosity fixed at $\mu = 1 \cdot 10^{-4}$, $5 \cdot 10^{-5}$, and $1 \cdot 10^{-5}$ for the varying simulations.




Figure 11: Contours of density for the triple-point shock interaction problem at varying times computed using a \mathbb{P}_3 approximation on a $N = 2800 \times 1200$ structured quadrilateral mesh.




Figure 12: Volume-averaged kinetic energy for the Rayleigh–Taylor instability problem computed using a \mathbb{P}_3 approximation on a $N = 32 \times 32 \times 64$ structured hexahedral mesh with varying values of the dynamic viscosity.

The problem was solved using a \mathbb{P}_3 approximation on the given mesh over the time range $t \in [0, 5]$. The quantity of comparison for the flow was the volume-averaged kinetic energy, computed as

$$k = \frac{1}{2} \rho \mathbf{V} \cdot \mathbf{V}. \quad (40)$$

A comparison of the kinetic energy profiles with respect to time for the varying values of the dynamic viscosity is shown in Fig. 12. The instabilities driven by the gravitational field introduce kinetic energy into the initially static flow field. This forcing is eventually balanced by the viscous dissipation of the turbulent mixing state at $t \approx 4.5$, after which viscosity overcomes and starts driving down the kinetic energy of the system. It can be seen that with decreasing values of the dynamic viscosity, the kinetic energy peak in the flow increases. Given the identical mesh used for the simulations, this indicates that the dissipation in the flow is primarily introduced through the physical viscosity, not the numerical stabilisation approach, which indicates that the proposed method may be a promising approach for accurately simulating complex multi-species turbulent flows.




Figure 13: Volume rendering of the density field for the Rayleigh–Taylor instability problem at varying times computed using a \mathbb{P}_3 approximation on a $N = 32 \times 32 \times 64$ structured hexahedral mesh with $\mu = 1 \cdot 10^{-5}$.

A visualisation of the flow through a volume rendering of the density field is also shown in Fig. 13 for the case of $\mu = 1 \cdot 10^{-5}$. The canonical flow structures of the descending “fingers” and ascending “pillars” of the Rayleigh–Taylor instability can be clearly seen in the early stages of the flow. The breakdown of these flow structures can then be observed at $t = 4$, after which a chaotic turbulent mixing state develops at $t = 5$. Eventually, the viscous dissipation in the system is expected to gradually bring the flow to a rest over longer simulation times.

5.2.2. Taylor–Green vortex

The final evaluation of the proposed method was performed on the Taylor–Green vortex, a canonical fluid dynamics problem for studying vortex dynamics and turbulent transition and decay. The problem consists of a laminar initial flow field which transitions to turbulence at Reynolds numbers of approximately 500 and above, which makes it an ideal case for evaluating the effects of the proposed scheme on predicting small-scale turbulent flow structures, particularly with respect to the numerical dissipation introduced by the scheme to stabilise the solution around species interfaces. To comprehensively study this, we consider the case of a two-species flow with identical specific heat constants and transport coefficients for both species which should, in theory, exhibit identical flow features to the single-species case. The resulting differences between the two cases can then be attributed to effects of the proposed numerical stabilisation method.

For this problem, the domain is taken as $\Omega = [0, 2\pi]^3$, and the initial conditions are defined as

$$\rho = 1, \quad (41a)$$

$$u = \sin x \cos y \cos z, \quad (41b)$$

Table 5

Overview of the various problem setups for the Taylor–Green vortex.

Test	Species (n)	α_0	N
A-coarse	1	1	32^3
A-fine	1	1	40^3
B-coarse	2	Eq. (42)	32^3
B-fine	2	Eq. (42)	40^3
C-coarse	2	Eq. (43)	32^3
C-fine	2	Eq. (43)	40^3

$$v = -\cos x \sin y \cos z, \quad (41c)$$

$$w = 0, \quad (41d)$$

$$P = P_0 + \frac{1}{16} (\cos(2x) + \cos(2y)) (\cos(2z) + 2), \quad (41e)$$

where the ambient pressure was set as $P_0 = 1/(\gamma M^2)$ for a specific heat ratio of $\gamma = 1.4$ and a Mach number of $M = 0.08$. The specific heat capacities of both fluid were set as $c_p = 1.005$ and $c_v = 0.718$. The Reynolds number and Prandtl number were set as $Re = 1600$ and $Pr = 0.71$, respectively. At these operating conditions, the Taylor–Green vortex is an extensively studied case [9].

Three variations of the problem were considered. The first variation, denoted by the label *A*, consists of the standard single-species description of the problem, where $\rho = \rho_0$. The second variation, denoted by the label *B*, consists of a two-species flow with an initially smooth species interface, given as

$$\alpha_0 = \frac{1}{2} + \frac{1}{16} (\cos(x) + 1)(\cos(y) + 1)(\cos(z) + 1). \quad (42)$$

The third variation, denoted by the label *C*, consists of a two-species flow with an initially sharp species interface, given as

$$\alpha_0 = \frac{1}{2} + \frac{1}{4} \tanh(\beta(x - \pi)) \tanh(\beta(y - \pi)) \tanh(\beta(z - \pi)), \quad (43)$$

where $\beta = 1000$. For each variation of the problem, both a coarse mesh, consisting of $N = 32^3$ elements, and a fine mesh, consisting of $N = 40^3$ elements, were considered. An overview of these problem setups is presented in Table 5.

The metric of interest in the flow was the dissipation rate measured by the enstrophy, defined as

$$\varepsilon = \frac{2\mu}{(2\pi)^3} \int_{\Omega} \frac{1}{2} \rho(\boldsymbol{\omega} \cdot \boldsymbol{\omega}) \, d\mathbf{x}, \quad (44)$$

where $\boldsymbol{\omega} = \nabla \times \mathbf{V}$ is the vorticity. For incompressible flows, this quantity is equal to the kinetic energy dissipation rate, but for compressible flows, it differs from the kinetic energy dissipation rate by the contribution of the deviatoric strain and pressure dilatation to the dissipation. As this functional is based on the vorticity in the flow field, an underprediction in the enstrophy indicates that the generation of small-scale flow features has been suppressed by numerical dissipation. The problem was solved using a \mathbb{P}_3 approximation, and the dissipation measured by enstrophy for the various problem setups on the coarse and fine mesh is presented in Fig. 14 in comparison to the direct numerical simulation (DNS) results of van Rees et al. [34] (as presented in Hillewaert [23]). It can be seen that the single-species results (*A*) show relatively good agreement with the DNS results in terms of enstrophy for the coarse mesh and show convergence with increasing resolution. When the flow field was described using a two-species approach (*B* and *C*), the measured enstrophy was slightly underpredicted with respect to the single-species approach. However, it was observed that these differences diminished with increasing resolution, with the case of the initially sharp interface (*C*) showing relatively good agreement with the single-species case. Notably, the case of the initially smooth interface (*B*) underpredicted the enstrophy in the flow more than the initially sharp interface, indicating that more numerical dissipation was introduced for a smoothly varying interface than a sharp interface. While this may seem counter-intuitive at first as it is expected that the numerical dissipation introduced by the proposed approach is primarily for the purpose of stabilising the

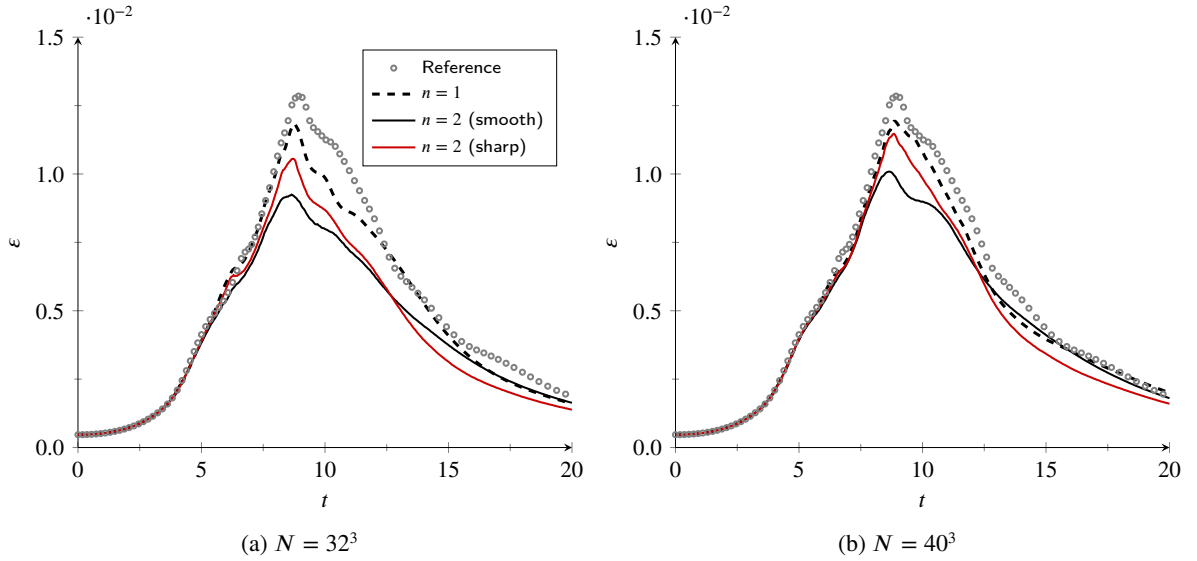


Figure 14: Dissipation measured by enstrophy for the Taylor–Green vortex computed using a \mathbb{P}_3 approximation on a $N = 32^3$ (left) and $N = 40^3$ (right) structured hexahedral mesh. Results shown for problem setups with single-species (black, dashed), multi-species with smooth phase boundaries (black, solid), and multi-species with sharp phase boundaries (red, solid). DNS results of van Rees et al. [34] (as presented in Hillewaert [23]

) shown for reference.

scheme in the vicinity of discontinuities, these observations may be linked to aliasing errors associated with the nonlinear description of the interface. It has been shown that the enforcement of entropy constraints tends to mitigate aliasing errors [13] through the addition of a suitable amount of numerical dissipation, which may explain the observed underprediction of the enstrophy. Regardless, the results indicate that the proposed approach can accurately resolve complex flow phenomena such as transition to turbulence for multi-species flows to a similar degree of accuracy that can be obtained by high-order DSEM for single-species flows.

6. Conclusions

In this work, we present a novel positivity-preserving numerical stabilisation approach for the conservative compressible multi-species flow equations solved via high-order discontinuous spectral element methods. The stabilisation method uses the adaptive nonlinear filtering approach of Dzanic and Witherden [14], with constraints applied to species density, total density, and pressure. It was found that the entropy constraints proposed in Dzanic and Witherden [14] led to excessive dissipation around species interfaces, substantially degrading the accuracy of the underlying high-order scheme. An improved stabilisation framework was proposed in this work where the entropy component of the constraints was adaptively enabled or disabled based on the presence of discontinuities in the pressure field, detected using a sensor which utilises the convergence properties of high-order DSEM. The efficacy of the proposed method was demonstrated in numerical experiments on the multi-species Euler and Navier–Stokes equations computed on structured and unstructured grids. It was shown that the scheme can recover the high-order accuracy of the underlying DSEM for smooth solutions while ensuring robustness in the vicinity of shock waves and contact discontinuities, with promising results in the ability to predict nonlinear flow phenomena such as flow instabilities and shock-vortex interactions. Furthermore, the proposed modification of the entropy constraints showed significant improvements in the ability of the method to accurately resolve species interfaces. Utilising the ability of the proposed framework to robustly and accurately resolve multi-species fluid flows, future work will focus on extending the approach to nonlinear equations of state and reacting multi-species flows.

Acknowledgements

WT would like to thank Robert Manson-Sawko for his insightful discussions. WT also acknowledges the support of UKRI through the grant MR/T041862/1 and the IBM Research Cognitive Computing Cluster service for providing resources that have contributed to the research results reported within this paper. TD would like to acknowledge the computational resources provided by the Princeton Institute for Computational Science and Engineering.

References

- [1] Abgrall, R., 1991. An extension of Roe's upwind scheme to algebraic equilibrium real gas models. *Computers & Fluids* 19, 171–182. doi:10.1016/0045-7930(91)90032-d.
- [2] Abgrall, R., 1996. How to prevent pressure oscillations in multicomponent flow calculations: A quasi conservative approach. *Journal of Computational Physics* 125, 150–160. doi:10.1006/jcph.1996.0085.
- [3] Abgrall, R., Karni, S., 2001. Computations of compressible multifluids. *Journal of Computational Physics* 169, 594–623. doi:10.1006/jcph.2000.6685.
- [4] Ansari, M., Daramizadeh, A., 2013. Numerical simulation of compressible two-phase flow using a diffuse interface method. *International Journal of Heat and Fluid Flow* 42, 209–223. doi:10.1016/j.ijheatfluidflow.2013.02.003.
- [5] Arrhenius, S., 1887. Über die innere reibung verdünnter wässriger lösungen. *Zeitschrift für Physikalische Chemie* 1U, 285–298. doi:10.1515/zpch-1887-0133.
- [6] Baer, M., Nunziato, J., 1986. A two-phase mixture theory for the deflagration-to-detonation transition (DDT) in reactive granular materials. *International Journal of Multiphase Flow* 12, 861–889. doi:10.1016/0301-9322(86)90033-9.
- [7] Bassi, F., Rebay, S., 2000. A high order discontinuous Galerkin method for compressible turbulent flows, in: *Lecture Notes in Computational Science and Engineering*. Springer Berlin Heidelberg, pp. 77–88. doi:10.1007/978-3-642-59721-3_4.
- [8] Billet, G., Ryan, J., 2011. A Runge–Kutta discontinuous Galerkin approach to solve reactive flows: The hyperbolic operator. *Journal of Computational Physics* 230, 1064–1083. doi:10.1016/j.jcp.2010.10.025.
- [9] Brachet, M.E., Meiron, D.I., Orszag, S.A., Nickel, B.G., Morf, R.H., Frisch, U., 1983. Small-scale structure of the Taylor–Green vortex. *Journal of Fluid Mechanics* 130, 411. doi:10.1017/s0022112083001159.
- [10] Chen, T., Shu, C.W., 2017. Entropy stable high order discontinuous Galerkin methods with suitable quadrature rules for hyperbolic conservation laws. *Journal of Computational Physics* 345, 427–461. doi:10.1016/j.jcp.2017.05.025.
- [11] Dowell, M., Jarratt, P., 1971. A modified regula falsi method for computing the root of an equation. *BIT* 11, 168–174. doi:10.1007/bf01934364.
- [12] Dunstan, T.D., Swaminathan, N., Bray, K.N.C., Cant, R.S., 2010. Geometrical properties and turbulent flame speed measurements in stationary premixed v-flames using direct numerical simulation. *Flow, Turbulence and Combustion* 87, 237–259. doi:10.1007/s10494-010-9284-1.
- [13] Dzanic, T., Trojak, W., Witherden, F.D., 2023. On the anti-aliasing properties of entropy filtering for under-resolved turbulent flows. *arXiv:arXiv:2302.13359*.
- [14] Dzanic, T., Witherden, F., 2022. Positivity-preserving entropy-based adaptive filtering for discontinuous spectral element methods. *Journal of Computational Physics* 468, 111501. doi:10.1016/j.jcp.2022.111501.
- [15] Dzanic, T., Witherden, F., 2023. Positivity-preserving entropy filtering for the ideal magnetohydrodynamics equations. *Computers & Fluids* 266, 106056. doi:10.1016/j.compfluid.2023.106056.
- [16] Galera, S., Maire, P.H., Breil, J., 2010. A two-dimensional unstructured cell-centered multi-material ALE scheme using VOF interface reconstruction. *Journal of Computational Physics* 229, 5755–5787. doi:10.1016/j.jcp.2010.04.019.
- [17] Gao, M., Appel, D., Beck, A., Munz, C.D., 2023. A high-order fluid–structure interaction framework with application to shock-wave/turbulent boundary-layer interaction over an elastic panel. *Journal of Fluids and Structures* 121, 103950. doi:10.1016/j.jfluidstructs.2023.103950.
- [18] Godunov, S.K., 1959. A difference method for numerical calculation of discontinuous solutions of the equations of hydrodynamics. *Matematicheskii Sbornik* 89, 271–306.
- [19] Gouasmi, A., Duraisamy, K., Murman, S.M., Tadmor, E., 2020. A minimum entropy principle in the compressible multicomponent Euler equations. *ESAIM: Mathematical Modelling and Numerical Analysis* 54, 373–389. doi:10.1051/m2an/2019070.
- [20] Grunberg, L., Nissan, A.H., 1949. Mixture law for viscosity. *Nature* 164, 799–800. doi:10.1038/164799b0.
- [21] Haas, J.F., Sturtevant, B., 1987. Interaction of weak shock waves with cylindrical and spherical gas inhomogeneities. *Journal of Fluid Mechanics* 181, 41. doi:10.1017/s0022112087002003.
- [22] Hesthaven, J.S., Warburton, T., 2008. *Nodal Discontinuous Galerkin Methods*. Springer New York. doi:10.1007/978-0-387-72067-8.
- [23] Hillewaert, K., 2012. Direct Numerical Simulation of the Taylor–Green Vortex at $Re = 1600$, in: *1st International Workshop on High-Order CFD Methods*, American Institute of Aeronautics and Astronautics.
- [24] Huynh, H.T., 2007. A flux reconstruction approach to high-order schemes including discontinuous Galerkin methods, in: *18th AIAA Computational Fluid Dynamics Conference*, American Institute of Aeronautics and Astronautics. doi:10.2514/6.2007-4079.
- [25] Jenny, P., Müller, B., Thomann, H., 1997. Correction of conservative Euler solvers for gas mixtures. *Journal of Computational Physics* 132, 91–107. doi:10.1006/jcph.1996.5625.
- [26] Johnson, R.F., Kercher, A.D., 2020. A conservative discontinuous Galerkin discretization for the chemically reacting Navier–Stokes equations. *Journal of Computational Physics* 423, 109826. doi:10.1016/j.jcp.2020.109826.
- [27] Kolev, T., Rieben, R., 2009. A tensor artificial viscosity using a finite element approach. *Journal of Computational Physics* 228, 8336–8366. doi:10.1016/j.jcp.2009.08.010.

- [28] Krivodonova, L., Xin, J., Remacle, J.F., Chevaugeon, N., Flaherty, J., 2004. Shock detection and limiting with discontinuous Galerkin methods for hyperbolic conservation laws. *Applied Numerical Mathematics* 48, 323–338. doi:10.1016/j.apnum.2003.11.002.
- [29] Kundu, A., 2021. Numerical simulation of a shock–helium bubble interaction. *Shock Waves* 31, 19–30. doi:10.1007/s00193-021-00996-x.
- [30] Kurnia, J.C., Sasmito, A.P., Mujumdar, A.S., 2014. CFD simulation of methane dispersion and innovative methane management in underground mining faces. *Applied Mathematical Modelling* 38, 3467–3484. doi:10.1016/j.apm.2013.11.067.
- [31] Michoski, C., Dawson, C., Kubatko, E.J., Wirasat, D., Brus, S., Westerink, J.J., 2015. A comparison of artificial viscosity, limiters, and filters, for high order discontinuous Galerkin solutions in nonlinear settings. *Journal of Scientific Computing* 66, 406–434. doi:10.1007/s10915-015-0027-2.
- [32] Peyvan, A., Shukla, K., Chan, J., Karniadakis, G., 2023. High-order methods for hypersonic flows with strong shocks and real chemistry. *Journal of Computational Physics* 490, 112310. doi:10.1016/j.jcp.2023.112310.
- [33] Quirk, J.J., Karni, S., 1996. On the dynamics of a shock-bubble interaction. *Journal of Fluid Mechanics* 318, 129. doi:10.1017/s0022112096007069.
- [34] van Rees, W.M., Leonard, A., Pullin, D., Koumoutsakos, P., 2011. A comparison of vortex and pseudo-spectral methods for the simulation of periodic vortical flows at high Reynolds numbers. *Journal of Computational Physics* 230, 2794–2805. doi:10.1016/j.jcp.2010.11.031.
- [35] Roe, P., 1981. Approximate Riemann solvers, parameter vectors, and difference schemes. *Journal of Computational Physics* 43, 357–372. doi:10.1016/0021-9991(81)90128-5.
- [36] Rusanov, V., 1962. The calculation of the interaction of non-stationary shock waves and obstacles. *USSR Computational Mathematics and Mathematical Physics* 1, 304–320. doi:10.1016/0041-5553(62)90062-9.
- [37] Sharan, N., Bellan, J.R., 2021. Direct numerical simulation of high-pressure free jets, in: *AIAA Scitech 2021 Forum*, American Institute of Aeronautics and Astronautics. doi:10.2514/6.2021-0550.
- [38] Shu, C.W., 1998. Essentially non-oscillatory and weighted essentially non-oscillatory schemes for hyperbolic conservation laws, in: *Lecture Notes in Mathematics*. Springer Berlin Heidelberg, pp. 325–432. doi:10.1007/bfb0096355.
- [39] Tadmor, E., 1986. A minimum entropy principle in the gas dynamics equations. *Applied Numerical Mathematics* 2, 211–219. doi:10.1016/0168-9274(86)90029-2.
- [40] Ton, V.T., 1996. Improved shock-capturing methods for multicomponent and reacting flows. *Journal of Computational Physics* 128, 237–253. doi:10.1006/jcph.1996.0206.
- [41] Tonicello, N., Ihme, M., 2023. A high-order diffused-interface approach for fully compressible two-phase flow simulations using a discontinuous Galerkin method doi:10.2514/6.2023-2154.vid.
- [42] Toor, H.L., 1969. Turbulent mixing of two species with and without chemical reactions. *Industrial & Engineering Chemistry Fundamentals* 8, 655–659. doi:10.1021/i160032a009.
- [43] Toro, E.F., 2009. The HLL and HLLC Riemann solvers, in: *Riemann Solvers and Numerical Methods for Fluid Dynamics*. Springer Berlin Heidelberg, pp. 315–344. doi:10.1007/b79761_10.
- [44] Wilke, C.R., 1950. A viscosity equation for gas mixtures. *The Journal of Chemical Physics* 18, 517–519. doi:10.1063/1.1747673.
- [45] Witherden, F., Farrington, A., Vincent, P., 2014. PyFR: An open source framework for solving advection–diffusion type problems on streaming architectures using the flux reconstruction approach. *Computer Physics Communications* 185, 3028–3040. doi:10.1016/j.cpc.2014.07.011.
- [46] Zhang, X., Shu, C.W., 2010. On maximum-principle-satisfying high order schemes for scalar conservation laws. *Journal of Computational Physics* 229, 3091–3120. doi:10.1016/j.jcp.2009.12.030.
- [47] Zhang, X., Shu, C.W., 2011. On positivity-preserving high order discontinuous Galerkin schemes for compressible euler equations on rectangular meshes. *Journal of Computational Physics* 229, 8917–8934. doi:10.1016/j.jcp.2010.08.016.
- [48] Zhang, X., Xia, Y., Shu, C.W., 2011. Maximum-principle-satisfying and positivity-preserving high order discontinuous Galerkin schemes for conservation laws on triangular meshes. *Journal of Scientific Computing* 50, 29–62. doi:10.1007/s10915-011-9472-8.