Implicit High-Order Flux Reconstruction Solver for High-Speed Compressible Flows

Ray Vandenhoeck^{a,*}, Andrea Lani^b

^a*Department of Mechanical Engineering, KU Leuven, Celestijnenlaan 300, 3001, Leuven, Belgium;*

Aeronautics & Aerospace Department, Von Karman Institute for Fluid Dynamics, Waterloosesteenweg 72, 1640 Sint-Genesius-Rode, Belgium

^b*Centre for Mathematical Plasma Astrophysics, KU Leuven, Celestijnenlaan 200B, 3001, Leuven, Belgium;*

Aeronautics & Aerospace Department, Von Karman Institute for Fluid Dynamics, Waterloosesteenweg 72, 1640 Sint-Genesius-Rode, Belgium

Abstract

The present paper addresses the development and implementation of the first high-order Flux Reconstruction (FR) solver for high-speed flows within the open-source COOLFluiD (Computational Object-Oriented Libraries for Fluid Dynamics) platform. The resulting solver is fully implicit and able to simulate compressible flow problems governed by either the Euler or the Navier-Stokes equations in two and three dimensions. Furthermore, it can run in parallel on multiple CPU-cores and is designed to handle unstructured grids consisting of both straight and curved edged quadrilateral or hexahedral elements. While most of the implementation relies on state-of-the-art FR algorithms, an improved and more case-independent shock capturing scheme has been developed in order to tackle the first viscous hypersonic simulations using the FR method. Extensive verification of the FR solver has been performed through the use of reproducible benchmark test cases with flow speeds ranging from subsonic to hypersonic, up to Mach 17.6. The obtained results have been favorably compared to those available in literature. Furthermore, so-called “super-accuracy” is retrieved for certain cases when solving the Euler equations. The strengths of the FR solver in terms of computational accuracy per degree of freedom are also illustrated and its computation cost per degree of freedom is assessed for different orders. Finally, the influence of the characterizing parameters of the FR method as well as the influence of the novel shock capturing scheme on the accuracy of the developed solver is discussed.

Keywords: High-Order Computational Fluid Dynamics, Flux Reconstruction, Correction Procedure via Reconstruction, Hypersonic Flows

Introduction

The field of Computational Fluid Dynamics (CFD) has been progressing significantly in recent years as a result of the need for increasingly complex simulations in a multitude of fields such as aerospace, automotive, wind engineering, atmospheric flows, etc. for a broad range of industrial and scientific applications. Depending on the spatial discretization of the Partial Differential Equations (PDEs) describing flow dynamics, the numerical methods developed in this field can be categorized into Finite Volume (FV), Finite Difference (FD), and Finite Element (FE) schemes. For each family of schemes, commercial codes used in the industry have largely been based on low-order methods which provide a second-order accuracy in space. Despite being generally robust, intuitive, and reliable, these codes have proven to be insufficiently accurate when applied to complex problems as encountered in the domains of turbulence, aero-acoustics, plasma flows, etc.

High-order methods can lead to considerably more accurate solutions for a comparable computational cost, i.e. for approximately the same number of degrees of freedom. Besides being computationally more efficient in terms of order of accuracy per degree of freedom, they also provide a lower numerical dissipation [1]. However, high-order methods are more complex to implement as compared to low-order methods and are often less robust. The fact that

*Corresponding author

Email addresses: ray.vandenhoeck@kuleuven.be (Ray Vandenhoeck), andrea.lani@kuleuven.be (Andrea Lani)

robust high-order mesh generators are not readily available has also limited their use for industrial applications [2]. Nonetheless, high-order methods have drawn considerable attention among researchers over the past couple of years due to their higher accuracy and their capability to deal with complex geometries [3]. Furthermore, the generation of high-order curved meshes has become increasingly supported by mesh generators such as the open-source Gmsh [4]. This has resulted in several successful attempts to develop high-order extensions of the existing FV, FD and FE schemes.

FD methods have been extended to higher orders in a straightforward manner by widening the stencil. In this regard, Compact Finite Difference [5] schemes have proven to provide higher orders of accuracy. Nonetheless, all FD methods exhibit the same weakness: they are incapable to readily deal with complex geometries, unless utilized within an overset framework which complicates mesh generation [6, 7, 8]. FV methods, on the other hand, do not have this shortcoming. As a consequence, several high-order extensions have been developed, such as the k -Exact method [9, 10, 11], the Essential Non-Oscillatory (ENO) [12, 13, 14], and Weighted Essential Non-Oscillatory (WENO) [15, 16, 17] schemes. These methods do not generally use a compact stencil and their utility for capturing complex fluid flow phenomena remains limited as the computational complexity rapidly increases at higher orders.

Due to the drawbacks of the high-order FV and FD extensions, several FE type methods that can attain arbitrarily high orders of accuracy on unstructured grids have been developed. In this respect, the Discontinuous Galerkin (DG) method was first proposed in 1973 by Reed and Hill [18]. Various extensions to the DG method have been proposed, such as Local Discontinuous Galerkin (LDG) [19], Compact Discontinuous Galerkin (CDG) [20] and Hybrid Discontinuous Galerkin (HDG) [21].

Among the various developed FE type methods, the Flux Reconstruction (FR) or Correction Procedure via Reconstruction (CPR) formulation [22, 23] is one of the most recent and promising family of schemes. The FR method was first proposed by Huynh in 2007 [22]. This approach provides a unifying framework for several existing high-order schemes, such as Spectral Difference (SD) [24] and DG, while being simpler and more computationally efficient than the original versions [25]. The relation between FR and DG has been studied thoroughly by Allaneau and Jameson [26], De Grazia et al. [27], and Zwanenburg et al. [28]. Furthermore, the FR approach is naturally adaptable to HPC architectures such as Graphical Processing Units (GPUs). Examples of massively parallel GPU implementations of the FR approach are given by Manuel et al. [29], Vincent et al. [30], and Witherden et al. [31]. The FR methods have been extended to work on any element types, including simplex element types in 2D and 3D [32, 33], and are thus suitable to handle arbitrarily complex geometries.

In 2011, Vincent, Castonguay, and Jameson proposed a class of FR schemes for linear advection problems in 1D, called Vincent-Castonguay-Jameson-Huynh (VCJH) schemes [34]. This method is characterized by a single parameter that determines the analytical form of the correction functions. Vincent, Castonguay, and Jameson proved the stability of the 1D VCJH schemes for linear advection using a similar energy method to the one Jameson used in [35] to prove the stability of an SD scheme for linear advection. As such, VCJH schemes are also referred to as Energy Stable Flux Reconstruction schemes (ESFR). Castonguay, Williams, Vincent, Lopez, and Jameson developed a multi-GPU enabled FR solver using the VCJH scheme, which is presented in [36].

The present paper treats a novel FR solver that is fully implicit and able to handle hypersonic viscous test cases. To the authors' knowledge, this is the first FR solver for hypersonic compressible flows. A novel shock capturing scheme for FR is introduced. The solver was implemented within COOLFluiD¹, a world-class open-source framework for multi-physics modeling and simulations where different numerical techniques, physical models, post-processing algorithms can coexist and work together [37, 38, 39].

Most shock capturing schemes for high-order FE-type methods that are currently being studied, are either adapted to explicit time stepping methods or unable to handle hypersonic viscous flows. This imposes a severe limit to the Courant-Friedrichs-Lewy (CFL) number and gives rise to a need for a large amount of iterations and computational time in order to reach convergence. In order to alleviate this issue, the present shock capturing scheme was successfully implemented for both explicit and implicit time stepping methods and verified in hypersonic test cases.

After giving an overview of the governing equations and the FR formulation in respectively section 1 and 2, the proposed shock capturing scheme is presented in section 3. Subsequently the implementation of the solver is discussed in section 4. Section 5 presents subsonic and hypersonic verification test cases both in inviscid and viscous regimes. Finally, section 6 presents an analysis of the performance of the solver.

¹<https://github.com/andrealani/COOLFluiD/wiki>

1. Governing Equations

Compressible flows are described by the Navier-Stokes equations. This set of equations is derived from the basic conservation laws: conservation of mass, momentum and energy where the fluid is assumed to be a continuum. This results in the following system of governing equations for three-dimensional flows:

$$\text{continuity: } \frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) = 0, \quad (1a)$$

$$\text{x-momentum: } \frac{\partial(\rho v_x)}{\partial t} + \nabla \cdot (\rho v_x \mathbf{v}) = \frac{\partial(\sigma_{xx} - p)}{\partial x} + \frac{\partial \tau_{yx}}{\partial y} + \frac{\partial \tau_{zx}}{\partial z}, \quad (1b)$$

$$\text{y-momentum: } \frac{\partial(\rho v_y)}{\partial t} + \nabla \cdot (\rho v_y \mathbf{v}) = \frac{\partial \tau_{xy}}{\partial x} + \frac{\partial(\sigma_{yy} - p)}{\partial y} + \frac{\partial \tau_{zy}}{\partial z}, \quad (1c)$$

$$\text{z-momentum: } \frac{\partial(\rho v_z)}{\partial t} + \nabla \cdot (\rho v_z \mathbf{v}) = \frac{\partial \tau_{xz}}{\partial x} + \frac{\partial \tau_{yz}}{\partial y} + \frac{\partial(\sigma_{zz} - p)}{\partial z}, \quad (1d)$$

$$\begin{aligned} \text{energy: } \frac{\partial(\rho e_t)}{\partial t} + \nabla \cdot (\rho e_t \mathbf{v}) = & -\nabla \cdot (p \mathbf{v}) + \left(\frac{\partial(v_x \sigma_{xx})}{\partial x} + \frac{\partial(v_x \tau_{yx})}{\partial y} \right. \\ & + \frac{\partial(v_x \tau_{zx})}{\partial z} + \frac{\partial(v_y \tau_{xy})}{\partial x} + \frac{\partial(v_y \sigma_{yy})}{\partial y} + \frac{\partial(v_y \tau_{zy})}{\partial z} + \frac{\partial(v_z \tau_{xz})}{\partial x} \\ & \left. + \frac{\partial(v_z \tau_{yz})}{\partial y} + \frac{\partial(v_z \sigma_{zz})}{\partial z} \right) - \nabla \cdot \boldsymbol{\varphi}, \end{aligned} \quad (1e)$$

where σ_{ii} denotes the normal viscous stress in the i -direction, and τ_{ij} denotes the shear viscous stress in the j -direction on a plane normal to the i -direction. ρ denotes the density, p the pressure and $\mathbf{v}$ the velocity vector. The heat flux vector $\boldsymbol{\varphi}$ is linearly proportional to the local temperature gradient as imposed by Fourier's law of heat conduction:

$$\boldsymbol{\varphi} = -\kappa \nabla T, \quad (2)$$

where κ represents the coefficient of thermal conductivity. Equations 1a through 1e are commonly known as the Navier-Stokes equations. When there is no viscosity and heat conduction present in the fluid, its flow is governed by the Euler equations. This implies that all viscous stresses σ_{ii} and τ_{ij} , as well as the thermal conductivity κ , are zero.

Assuming a Newtonian fluid, the viscous stresses are expressed as a function of the rates of deformation:

$$\sigma_{xx} = 2\mu \frac{\partial v_x}{\partial x} + \lambda \nabla \cdot \mathbf{v}, \quad \sigma_{yy} = 2\mu \frac{\partial v_y}{\partial y} + \lambda \nabla \cdot \mathbf{v}, \quad \sigma_{zz} = 2\mu \frac{\partial v_z}{\partial z} + \lambda \nabla \cdot \mathbf{v}, \quad (3)$$

$$\tau_{xy} = \tau_{yx} = \mu \left(\frac{\partial v_x}{\partial y} + \frac{\partial v_y}{\partial x} \right), \quad \tau_{xz} = \tau_{zx} = \mu \left(\frac{\partial v_x}{\partial z} + \frac{\partial v_z}{\partial x} \right), \quad \tau_{yz} = \tau_{zy} = \mu \left(\frac{\partial v_y}{\partial z} + \frac{\partial v_z}{\partial y} \right). \quad (4)$$

For gases, it is common practice to adopt Stokes' hypothesis, which states that $\lambda = -\frac{2}{3}\mu$. Furthermore, the ideal gas equation is used to close the system of equations.

The system of equations 1 can be rewritten in the compact form, called the conservative form:

$$\frac{\partial \mathbf{u}}{\partial t} + \nabla \cdot \mathbf{f}(\mathbf{u}, \nabla \mathbf{u}) = \frac{\partial \mathbf{u}}{\partial t} + \nabla \cdot \mathbf{f}_C(\mathbf{u}) - \nabla \cdot \mathbf{f}_D(\mathbf{u}, \nabla \mathbf{u}) = 0, \quad (5)$$

where $\mathbf{u} = (\rho, \rho v_x, \rho v_y, \rho v_z, \rho e_t)^T$ denotes the set of conserved variables and $\mathbf{f}(\mathbf{u})$ the set of total flux vectors. Consequently $\nabla \cdot \mathbf{f}(\mathbf{u})$ denotes the set of the divergences of the flux vectors. The flux vectors can be split up into a convective flux $\mathbf{f}_C(\mathbf{u})$, and a diffusive flux $\mathbf{f}_D(\mathbf{u}, \nabla \mathbf{u})$.

2. Flux Reconstruction Formulation

This section presents a concise review of the Flux Reconstruction method in one dimension equivalent to the method presented in [34, 36]. Consider solving the one-dimensional advection-diffusion problem on an arbitrary

domain Ω , given by:

$$\frac{\partial u}{\partial t} + \frac{\partial f}{\partial x} = 0 \quad \text{with} \quad x \in \Omega = [x_L, x_R] \quad \text{and} \quad f = f\left(u, \frac{\partial u}{\partial x}\right). \quad (6)$$

In equation 6, x represents the one dimensional coordinate, t is the time variable, $u(x, t)$ is a conservative variable, and f is the flux which is a scalar for the 1D case. The Euler and Navier-Stokes equations can be written as a system of conservation equations, equivalent to expression 6. The FR method is classified as a FE-type method. Hence, the spatial domain Ω is partitioned into a finite number N of non-overlapping, non-empty, open sub-domains Ω'_n such that:

$$\Omega'_n = \{x \mid x_n < x < x_{n+1}\} \text{ with } x_1 = x_L, \quad x_{N+1} = x_R. \quad (7)$$

The closed domain corresponding to the union of the open sub-domain Ω'_n and its boundary $\Gamma_n = \{x_n, x_{n+1}\}$ is called an element and is denoted by $\Omega_n = \Omega'_n \cup \Gamma_n$ such that:

$$\bigcup_{n=1}^N \Omega_n = \Omega \quad \text{with} \quad \bigcap_{n=1}^N \Omega_n = \emptyset. \quad (8)$$

The exact solution u of the conservation law 6 is approximated by a function u_n^δ within each element Ω_n . This function u_n^δ corresponds to a polynomial of degree P within Ω_n and is identically zero on $\Omega \setminus \Omega_n$. These polynomials are generally discontinuous across elements. u^δ is constructed through summation of the elemental solutions u_n^δ . In the same manner, the exact flux f is approximated by a polynomial of degree $P + 1$ within each sub-domain Ω_n , denoted by f_n^δ . The overall approximate flux f^δ in the entire domain Ω is given by the sum of the elemental approximate fluxes such that:

$$u^\delta(x, t) = \sum_{n=1}^N u_n^\delta(x, t) \quad \text{and} \quad f^\delta(x, t) = \sum_{n=1}^N f_n^\delta(x, t). \quad (9)$$

Each element Ω_n together with the approximate solution u_n^δ and flux f_n^δ inside it, is transformed to a standard reference element $\Omega_S = \{\xi \mid -1 \leq \xi \leq 1\}$. In this manner, the computations for all elements are done in the same reference domain Ω_S . Consequently, the same algorithm can easily be applied to all elements. The transformation between Ω_n and Ω_S is carried out by means of a mapping function $\Theta_n(\xi)$.

The solution u_n^δ within each sub-domain Ω_n can be obtained by solving the transformed conservation equation within the reference element Ω_S :

$$\frac{\partial \hat{u}^\delta}{\partial t} + \frac{\partial \hat{f}^\delta}{\partial \xi} = 0 \quad \text{with} \quad \xi \in \Omega_S, \quad (10)$$

with the transformed physical quantities given by:

$$\hat{u}^\delta = \hat{u}^\delta(\xi, t) = J_n u_n^\delta(\Theta_n(\xi), t) \quad \text{and} \quad \hat{f}^\delta = \hat{f}^\delta(\xi, t) = f_n^\delta(\Theta_n(\xi), t). \quad (11)$$

In equation 11, J_n represents the Jacobian of the mapping $\Theta_n(\xi)$.

The FR method for solving equation 10 within the reference element Ω_S consists of seven subsequent steps. In the first step, a specific form for the approximate solution $\hat{u}^\delta$ within Ω_S is defined. To this end, a set of $P + 1$ distinct solution points ξ_i are chosen within Ω_S at which the values of $\hat{u}^\delta$ are assumed to be known. The two end-points of the standard element Ω_S are referred to as the flux points.

The approximate solution $\hat{u}^\delta$ is defined as a degree P polynomial of the following form:

$$\hat{u}^\delta = \sum_{i=1}^{P+1} \hat{u}_i^\delta \ell_i(\xi), \quad (12)$$

where $\hat{u}_i^\delta$ is the value of $\hat{u}^\delta$ at the i -th solution point, and $\ell_i(\xi)$ is the 1D Lagrange polynomial associated with the i -th solution point.

The second step consists of determining a common interface solution $\hat{u}^{\delta I}$ at the boundaries of the standard element Ω_S , i.e. $\xi = \pm 1$, where the flux points of two neighboring elements coincide. In order to compute this interface value, the approximate solution $\hat{u}^\delta$ must first be evaluated in the flux points by means of equation 12. The approximate

solutions evaluated at the left and right boundaries are denoted as $\hat{u}_L^\delta = \hat{u}^\delta(-1)$ and $\hat{u}_R^\delta = \hat{u}^\delta(1)$. Following the second approach of Bassi and Rebay (BR2), $u^{\delta I}$ is then computed as the average of the two solutions in a flux point [40]. The common interface solutions corresponding to the left and right ends of Ω_n are denoted by $u_L^{\delta I}$ and $u_R^{\delta I}$. Figure 1 illustrates the process of computing the common interface values.

Figure 1: Computation of the common interface solutions $u_L^{\delta I}$ and $u_R^{\delta I}$ at the left and right boundary of Ω_n

The third step of the FR method consists of calculating a corrected solution gradient $\hat{q}^\delta$. The corrected solution gradient $\hat{q}^\delta$ is defined in terms of two correction functions $g_L(\xi)$ and $g_R(\xi)$ of degree $P + 1$, which approximate the zero function within the reference element Ω_S , and satisfy the following boundary conditions: [36]

$$g_L(-1) = 1, g_R(-1) = 0, g_L(1) = 0, g_R(1) = 1. \quad (13)$$

Furthermore, the correction function for the right boundary $g_R(\xi)$ is defined as the reflection of the correction function $g_L(\xi)$ with respect to the vertical axis: $g_L(\xi) = g_R(-\xi)$.

The corrected solution gradient $\hat{q}^\delta$ is then computed via the following expression:

$$\hat{q}^\delta = \frac{\partial \hat{u}^\delta}{\partial \xi} + (\hat{u}_L^{\delta I} - \hat{u}_L^\delta) \frac{dg_L}{d\xi} + (\hat{u}_R^{\delta I} - \hat{u}_R^\delta) \frac{dg_R}{d\xi}. \quad (14)$$

The VCJH scheme is used to define the exact polynomial expression of g_L and g_R , as found in [41].

The fourth step consists of calculating the approximate discontinuous flux $\hat{f}^{\delta D}(\xi)$ on the reference element Ω_S . To this end, the flux is computed at each solution point ξ_i . In general, the flux depends on the solution and the solution gradient. Consequently, the discontinuous flux is evaluated using the approximate solution $\hat{u}^\delta$, and the corrected solution gradient $\hat{q}^\delta$. The approximate discontinuous flux $\hat{f}^{\delta D}$ is then constructed as a polynomial of degree P in the same manner as $\hat{u}^\delta$ in equation 12:

$$\hat{f}_i^{\delta D} = f(\hat{u}_i^\delta / J_n, \hat{q}_i^\delta / J_n) = \sum_{i=1}^{P+1} \hat{f}_i^{\delta D} \mathcal{L}_i(\xi). \quad (15)$$

The fifth step of the FR framework involves computing the transformed common interface fluxes $\hat{f}^{\delta I}$ at the flux points of the standard element Ω_S . In order to compute these interface fluxes, it is necessary to first obtain the values of the approximate solution $\hat{u}^\delta$ and the corrected gradient $\hat{q}^\delta$ at the boundaries of the standard element by using equations 12 and 14. The common interface flux at the left and right ends of the standard element Ω_S is denoted as $\hat{f}_L^{\delta I}$ and $\hat{f}_R^{\delta I}$, respectively. The methodology for calculating the interface fluxes depends on the system of equations that needs to be solved. In the case of the Navier-Stokes equations, it is computed as the sum of a convective and a diffusive term. For the convective interface flux, an upwind biased approximate Riemann solver, like the Roe scheme [42], can be used. The diffusive interface flux is often calculated following the BR2 scheme, which uses the average approximate solution $((u^{\delta-} + u^{\delta+})/2)$ and the average corrected gradient $((q^{\delta-} + q^{\delta+})/2)$ as demonstrated by Bassi and Rebay [40] and applied to the FR approach in [36].

The flux must be continuous across element boundaries in order to obtain a conservative scheme. Therefore, the sixth stage of the FR framework consists of adding a correction flux function $\hat{f}^{\delta C}$ to the approximate discontinuous flux $\hat{f}^{\delta D}$ resulting in the continuous flux $\hat{f}^\delta$. The correction flux $\hat{f}^{\delta C}$ is a polynomial of degree $P + 1$ defined in such a manner that the continuous flux $\hat{f}^\delta$ equals the common interface flux in the flux points, while approximating the discontinuous flux $\hat{f}^{\delta D}$ within the standard element Ω_S . In order to satisfy the aforementioned requirements, $\hat{f}^{\delta C}$ is constructed by means of the degree $P + 1$ correction functions $h_L(\xi)$ and $h_R(\xi)$. These polynomials are equivalent to

the correction functions g_L and g_R used to define the corrected gradient $\hat{q}^\delta$ and have the same properties as expressed in 13.

The correction flux function $\hat{f}^{\delta C}$ takes the form:

$$\hat{f}^{\delta C} = (\hat{f}_L^{\delta I} - \hat{f}_L^{\delta D})h_L(\xi) + (\hat{f}_R^{\delta I} - \hat{f}_R^{\delta D})h_R(\xi), \quad (16)$$

where $\hat{f}_L^{\delta D} = \hat{f}^{\delta D}(-1)$ and $\hat{f}_R^{\delta D} = \hat{f}^{\delta D}(1)$ are the values of the approximate discontinuous flux at the left and right boundaries, evaluated from equation 15. The continuous flux $\hat{f}^\delta$ within the standard element Ω_S is then written as the sum of the approximated discontinuous flux and the correction flux:

$$\hat{f}^\delta = \hat{f}^{\delta D} + (\hat{f}_L^{\delta I} - \hat{f}_L^{\delta D})h_L + (\hat{f}_R^{\delta I} - \hat{f}_R^{\delta D})h_R. \quad (17)$$

A schematic of the flux correction procedure within the reference element Ω_S is given in figure 2.

Figure 2: Correction procedure $\hat{f}^{\delta D}$ such that the values of the total flux $\hat{f}^\delta$ at the boundaries are equal to the common interface fluxes $\hat{f}_L^{\delta I}$ and $\hat{f}_R^{\delta I}$

The final step consists of computing the divergence of the total flux $\hat{f}^\delta$ at each solution point ξ_i , which is trivial as the derivatives of both the Lagrange polynomials and correction functions are exactly known.

Finally, a time marching strategy such as a backward Euler method is applied to advance the approximate solution $\hat{u}^\delta$ in time using the following expression:

$$\frac{d\hat{u}_i^\delta}{dt} = -\frac{\partial \hat{f}^\delta}{\partial \xi}(\xi_i). \quad (18)$$

For quadrilateral and hexahedral elements, the extension to 2D and 3D is straightforward: it suffices to define a tensor product basis based on the one-dimensional Lagrange polynomials on the reference element, as shown in [22]. Subsequently the steps of the FR method are analogous to FR in one dimension.

3. Shock Capturing

Robust shock capturing is the main pacing item for high-order FE-type CFD methods. In the vicinity of discontinuities within the flow field, spurious oscillations appear due to the Gibbs phenomenon. This effect is more severe for higher orders and for stronger shocks. The oscillations near the shock generally cause numerical instabilities and negative pressures or densities. Several schemes have been investigated in [43, 44, 45, 46, 47] to avoid spurious oscillations around shock waves. However these schemes have not been successfully applied to viscous hypersonic flows and are generally only applicable to explicit time stepping. The present paper proposes a modified Localized Laplacian Artificial Viscosity (LLAV) scheme combined with a positivity preservation method in order to alleviate oscillations caused by the Gibbs phenomenon as well as harsh transient effects characterizing hypersonic flow calculations.

3.1. Localized Laplacian Artificial Viscosity

The concept of localized artificial viscosity is to add an artificial diffusive flux dependent on the local properties of the flow field in order to spread out the shock over one or more elements. In this way spurious oscillations around shocks are attenuated. Within the LLAV formulation, the artificial viscosity is of a Laplacian type, as proposed in [46, 48]. The local artificial viscosity is only activated in areas of the flow field where discontinuities are present. To this end, a smoothness detector is used, of which several are proposed in [46].

LLAV shows better convergence characteristics compared to traditional limiters developed for FR, such as Multi-Dimensional Limiting Process (MLP) [44], especially for steady-state cases. Furthermore LLAV is readily adaptable to implicit time-marching schemes. However, LLAV cannot guarantee positivity of the solution for hypersonic test cases and requires fine-tuning of the maximum amount of artificial viscosity added, to work properly. A modified LLAV scheme is proposed to alleviate these difficulties.

Within the LLAV framework, an extra flux is added to the system of equations:

$$\frac{\partial \mathbf{u}}{\partial t} = -\nabla \cdot \mathbf{f}_C + \nabla \cdot \mathbf{f}_D + \nabla \cdot \mathbf{f}_{AV}, \quad (19)$$

where $\mathbf{f}_C$ is the convective flux, $\mathbf{f}_D$ the diffusive flux and $\mathbf{f}_{AV}$ the flux due to the artificial viscosity. $\mathbf{f}_{AV}$ is constructed as follows:

$$\mathbf{f}_{AV} = \varepsilon \nabla \mathbf{u}, \quad (20)$$

where ε is the artificial viscosity. Within the FR framework, this means that the space discretization procedure, as presented in section 2, is also applied to a third flux. The gradients $\nabla \mathbf{u}$ are approximated using the corrected solution gradients $\mathbf{q}$ as computed in equation 14. As such, the FR procedure is applied to $\mathbf{f}_{AV} = \varepsilon \mathbf{q}$.

The value of the artificial viscosity ε is computed in such a way that it is only active in discontinuous regions:

$$\varepsilon = \begin{cases} 0 & \text{if } S < S_0 - \kappa, \\ \frac{\varepsilon_0}{2} \left(1 + \sin \frac{\pi(S - S_0)}{2\kappa} \right) & \text{if } S_0 - \kappa \leq S \leq S_0 + \kappa, \\ \varepsilon_0 & \text{if } S > S_0 + \kappa, \end{cases} \quad (21)$$

where S is the smoothness in the current element, S_0 the reference smoothness, κ a parameter to control the working spectrum of the artificial viscosity and ε_0 the maximum artificial viscosity.

The smoothness for a P -th order FR method can be computed as follows, using the smoothness detector S found in [46]:

$$S = \log_{10} \frac{\langle u_m - u_m^{P-1}, u_m - u_m^{P-1} \rangle}{\langle u_m, u_m \rangle}, \quad (22)$$

where u_m is the state variable that is monitored in order to detect discontinuities. For the Navier-Stokes equations, typically the density ρ or pressure p is used. u_m^{P-1} represents the monitored state variable projected on a $(P - 1)$ -th order polynomial. The operator $\langle \cdot, \cdot \rangle$ is the element-wise scalar product over the solution points. The reference smoothness is set equal to $S_0 = -3 \log_{10} P$ as suggested in [49]. The $(P - 1)$ -projected state variable can be computed by projecting the solution polynomial, which is expression in a Lagrange polynomial basis, on a modal basis using the Vandermonde matrix. Next, the sum over the first $P - 1$ polynomials is made to determine u_m^{P-1} .

A novel scheme to compute the maximum artificial viscosity is proposed. In order to match the flux dimensions, ε must be expressed in $\left[\frac{m^2}{s} \right]$, or a length-scale times a velocity-scale. Since a higher artificial viscosity is needed at higher speeds and in bigger elements (in order to spread out the shock over the element), ε_0 is set to be:

$$\varepsilon_0 = f(P)h|\lambda|_{max}, \quad (23)$$

where h is the characteristic length of the current element and $|\lambda|_{max}$ the maximum eigenvalue of the inviscid part of the system of equations. In order to take into account the order P of the FR method, following the approach of [46], $f(P)$ is set to:

$$f(P) = c(2 - \Delta\xi_{max}(P)), \quad (24)$$

where $\Delta\xi_{max}(P)$ is the subcell resolution, i.e. the largest distance between two solution points within the reference domain and c a constant. In order to avoid the need for fine-tuning the constant c , an approximate relation between this parameter and the number of cells over which a shock is spread is defined here after.

For strong shocks where viscosity is present, the following shock-related Reynolds number is of order of magnitude 1:

$$\text{Re}_s = \frac{\delta_s \Delta u \rho^*}{\mu^*} \approx 1, \quad (25)$$

where δ_s is the shock thickness, Δu the velocity difference over the shock, ρ^* and μ^* respectively the density and dynamic viscosity in the shock at $M = 1$. Consider the following assumptions:

$$\rho^* \sim \rho_\infty, \quad (26)$$

$$\mu^* \sim \varepsilon \rho_\infty. \quad (27)$$

Combining these two relations, the following scaling relation is found:

$$\varepsilon \rho_\infty \sim \delta_s \Delta u \rho_\infty. \quad (28)$$

Using the expression for the velocity drop Δu over a shock:

$$\varepsilon \sim \delta_s \frac{2}{\gamma + 1} \frac{M_\infty^2 - 1}{M_\infty^2}, \quad (29)$$

combining with expression 23:

$$c(2 - \Delta \xi_{max}) h |\lambda|_{max} \sim \delta_s \frac{M_\infty^2 - 1}{M_\infty^2}, \quad (30)$$

rearranging and setting $h^* = (2 - \Delta \xi_{max}) h$ the characteristic subcell length:

$$c \sim \frac{\delta_s}{h(u_\infty + \frac{u_\infty}{M_\infty})} \frac{M_\infty^2 - 1}{M_\infty^2}, \quad (31)$$

$$c \sim \frac{\delta_s}{h^*} \frac{M_\infty - 1}{M_\infty}, \quad (32)$$

introducing the characteristic shock thickness $\delta_s^* = \frac{\delta_s}{h^*}$. This variable is a measure for the number of characteristic element lengths over which the shock will span. Using this variable, the following relation is found:

$$c \sim \delta_s^* \frac{M_\infty - 1}{M_\infty}. \quad (33)$$

Using expression 33, the LLAV scheme can be characterized by a characteristic shock thickness δ_s^* which is a measure for the number of elements over which the shock is spread out. By calibrating c as:

$$c = \alpha \delta_s^* \frac{M_\infty - 1}{M_\infty} + \beta. \quad (34)$$

If there is no viscosity present, the characteristic shock thickness is zero. As a consequence, β can be set to zero. For the value of α , a calibrating test case can be used. As described in the section 5, the following value is found:

$$\alpha \simeq 0.027. \quad (35)$$

By using this scheme, only the parameters κ and δ_s^* need to be specified for the LLAV method instead of κ and c . For the spectral sensitivity κ a value ranging from 1 to 5 generally produces robust results, as proposed in [46]. This parameter still needs to be fine-tuned between those values for each separate test case. This remains a disadvantage of the present scheme. The characteristic shock thickness is a measure for the amount of elements over which the shock is spread. As a consequence, if $\delta_s^* < 1$, there is a risk of spurious oscillations, while if $\delta_s^* \gg 2$, the shock will be spread over more than two elements, and accuracy may be deteriorated due to superfluous artificial viscosity. The LLAV scheme based on these two parameters requires little fine-tuning: only for the parameter κ , and it generally produces robust results, independent of the characteristic length h of the elements near the shock, of the Mach number of the far field, etc. In this way the fine-tuning of the novel LLAV scheme defined by δ_s^* is much more straightforward than that of the LLAV scheme defined by the scaling parameter c .

3.2. Positivity Preservation Scheme

For the Navier-Stokes equations the conservative state variables are $\mathbf{u} = (\rho, \rho\mathbf{v}, \rho e_t)$, where $\mathbf{v}$ is the velocity and e_t the specific total energy. However not all states are admissible: density and pressure must remain positive:

$$\rho > 0 \quad \text{and} \quad p = (\gamma - 1) \left(\rho e_t - \frac{1}{2} \frac{\|\rho\mathbf{v}\|^2}{\rho} \right) > 0. \quad (36)$$

If the above conditions are not satisfied, a numerical blow-up of the solution will occur. A shock capturing method based on artificial viscosity cannot guarantee that equation 36 is satisfied. Especially in hypersonic test cases where very strong shocks and severe transients are present, artificial viscosity is inadequate in preserving positivity. In order to alleviate this, a positivity preservation scheme is proposed for FR. This scheme limits the state at each iteration such that each violation of equation 36 is corrected. The proposed scheme is similar to the one presented in [50] that was developed for the DG method, with the difference that now the limiting procedure must limit the states based on the evaluation of the approximate solution polynomials in the flux points.

The first step is to limit the density. To this end the average state $\bar{\mathbf{u}}$ on the each element is computed. Based on $\bar{\mathbf{u}}$, the average density $\bar{\rho}$ and pressure $p(\bar{\mathbf{u}})$ are known. A small number ϵ is defined to which ρ will be limited:

$$\epsilon = \min(10^{-13}, \bar{\rho}, p(\bar{\mathbf{u}})), \quad (37)$$

where $p(\bar{\mathbf{u}})$ denotes the pressure computed using the element average state. Next, ρ is recomputed in each solution point i , in order to obtain a new positive density $\tilde{\rho}_i$:

$$\tilde{\rho}_i = t_1(\rho_i - \bar{\rho}) + \bar{\rho} \quad \text{with} \quad t_1 = \min\left(\frac{\bar{\rho} - \epsilon}{\bar{\rho} - \rho_{min}}, 1\right). \quad (38)$$

ρ_{min} is computed in the following way:

$$\rho_{min} = \min_j(\rho_j), \quad (39)$$

where the subscript j refers to the j -th flux point. In this way, the density in each solution point of an element is limited such that all the densities in the flux points remain larger than or equal to ϵ .

The state $\tilde{\mathbf{u}} = (\tilde{\rho}, \rho\mathbf{v}, \rho e_t)$ obtained after limiting the density, is now used to limit the pressure: if in the j -th flux point $p_j < \epsilon$, a coefficient t_j is computed such that:

$$p(t_j(\tilde{\mathbf{u}}_j - \bar{\mathbf{u}}) + \bar{\mathbf{u}}) = \epsilon. \quad (40)$$

$\tilde{\mathbf{u}}_j$ represents the state after limiting the density in the j -th flux point. This can be computed by evaluating the approximate solution polynomial, as defined in equation 12, in the j -th flux point. The value of t_j can be obtained from equation 40 by solving a second-order polynomial equation. Note that the expression $0 < t_j < 1$ is used to select the correct solution of the second-order equation. Next, the whole state in each solution point i is limited using the minimum value of the t_j coefficients computed in the current element, obtaining $\tilde{\tilde{\mathbf{u}}}_i$:

$$\tilde{\tilde{\mathbf{u}}}_i = t_2(\tilde{\mathbf{u}}_i - \bar{\mathbf{u}}) + \bar{\mathbf{u}} \quad \text{with} \quad t_2 = \min_j(t_j). \quad (41)$$

This scheme guarantees that the density and pressure remain positive in the flux points. For a $P1$ extrapolation, the most negative value will always be attained inside a flux point rather than a solution point, since they are situated on the boundary of the element. For higher order it is possible that a negative state appears first in a solution point. As such, the same procedure as outlined above is also applied to the solution points.

3.3. Interaction between LLAV and positivity preservation

Generally for a high maximum artificial viscosity, which is accomplished by specifying a large δ_s^* , positivity is only enforced during the transient of the simulation. However if δ_s^* is chosen lower, typically $\delta_s^* \simeq 1$, the positivity preservation scheme remains active until convergence. Since the positivity preservation scheme is applied outside of the residual computation, it causes the residuals to stall at convergence: the positivity preservation scheme exactly

matches the contribution of the non-zero residuals to the solution at the next time step. As a consequence the solution remains unchanged, despite non-zero residuals. As proposed in [47], a solution for this is to monitor $\frac{u(t+\Delta t)-u(t)}{\Delta t}$ instead of the residual. This is a behavior typically found in limiter-type shock capturing methods for high-order FE schemes, such as Multi-Dimensional Limiting Process, presented in [47]. However, since the positivity preservation scheme aims at keeping the solver stable during the transient, it should be avoided that it remains active until convergence. δ_s^* can simply be increased in order to avoid this behavior. In the present solver implementation it is possible to do this “on the fly” during the simulation.

Furthermore, the positivity preservation scheme is applied outside of the residual computation process and imposes a severe limit to the CFL number in implicit time marching. This is due to the fact that the contribution of the positivity preservation scheme in the Jacobian of the implicit scheme cannot readily be evaluated. However, since the positivity preservation scheme should only be active during the transient, the severe CFL limit is only imposed as long as it remains active.

4. Implementation of the Solver

This section first presents an overview of the capabilities of the developed solver. Second, the general structure is described as well as the main interactions between various parts of the solver and the existing COOLFluid framework. The present FR solver consists of two main segments: the data structure and the FR algorithm itself. Both parts are successively discussed in detail.

The present FR solver is implemented as a plug-in of COOLFluid, a platform written in C++ that offers a component-based framework oriented towards complex multi-physics simulations. Herein, each numerical method or physical model is encapsulated into an independent dynamic module (or *plug-in* library) that can be loaded on demand by user-defined applications. Some of the main features of COOLFluid include: parallel solvers for compressible and incompressible flows based on multiple discretization techniques (e.g. FV, Residual Distribution schemes, Finite Element, Spectral Finite Differences) for unstructured meshes, interfaces to different linear systems solvers (e.g. PETSc, Trilinos, Pardiso), aerothermochemical models for flows and plasma [51, 52, 53, 54, 55, 56], MHD, coupling algorithms for multi-physics and multi-domain simulations, Arbitrary Lagrangian Eulerian methods and radiation transport algorithms based on Monte Carlo [37, 57]. More information on the development of this platform can be found in [58, 59].

The FR solver has been developed to solve the compressible Euler and Navier-Stokes equations in 2D and 3D. In COOLFluid, the algorithms that depend on the physics of a problem, such as the evaluation of fluxes, are decoupled from the numerical algorithms. This means that the FR solver has been developed to solve advection and advection-diffusion problems in general. In addition, all necessary physics-dependent algorithms for the Euler and Navier-Stokes equations were implemented. These are the boundary conditions that compute the ghost states and gradients as well as the algorithm that computes the characteristic speeds for the evaluation of the CFL number. Extending the functionality of the FR solver to different types of advection-diffusion problems only requires the addition of the above mentioned physics-dependent functions.

The FR method is fully determined by the solution and flux point distributions, the type of correction function and the interface flux scheme. The present FR solver can employ Gauss-Legendre, Gauss-Lobatto or equidistant points for both solution and flux points. These point distributions are defined in separate C++ classes within the solver. The user can specify which distribution must be used. As a consequence, it suffices to add an extra class without changing the existing solver, to add new distributions. The VCJH scheme for the correction functions has been implemented in the present solver for quadrilaterals and hexahedra. The following interface flux schemes have been implemented: Centered flux [60], Lax-Friedrichs flux [61], Roe flux [42], AUSM⁺ [62] and AUSM⁺-up flux [63]. These methods are used to compute the common interface fluxes. Furthermore, the FR method has been implemented from P1 up to P9.

Finally, the solver can either be executed serially or in parallel. For the serial execution, one CPU computes the update values of the whole mesh at each iteration. For the parallel execution, the spatial domain is partitioned between a number of CPUs specified by the user. Details of the partitioning procedure, which are out of the scope of the present paper, are a high-level feature of COOLFluid and are explained in [58, 64].

The developed FR solver consists of two major components: the FR data structure and the FR algorithm. The main function of the data structure is to create the data that defines the high-order unstructured mesh. This is com-

prised of geometric entities, namely elements and faces, how these elements are connected to each other, and their properties: solution points, flux points, base functions, etc. The FR algorithm subsequently uses this data to compute the numerical solution of a given test case in an iterative way.

Both parts of the FR solver interact with the existing COOLFluid framework, as schematically shown in figure 3. In the next sections, the two major parts of the solver are discussed in detail.

Figure 3: Overview of the major parts of the FR solver

4.1. Implementation of the Data Structure

The FR data structure allows the solver to loop over the geometric entities of the mesh. Geometric entities are either elements or faces. In addition, a mesh is specified by a list of nodes and states. For unstructured meshes there is no trivial way to access neighboring geometric entities. As a consequence, the creation of look-up tables, called connectivity tables, is an important function of the data structure. These tables store how geometric entities are linked to each other. The collection of a certain type of geometric entities and their relevant connectivity tables is referred to as a Topological Region Set (TRS). TRSs include the boundary as well as the internal part of the computational domain. The first function of the FR data structure is the creation of TRSs. Since the present FR solver builds the high-order TRSs by upgrading the order of the mesh data generated in the COOLFluid framework, the block “Mesh Upgrade” in figure 3 represents this functionality. The second function of the data structure is to store the geometric data and properties of the different element types that are needed for the FR algorithm.

4.1.1. Building the Topological Region Sets

The FR data structure creates the following topological region sets: a set of all elements, inner faces, boundary faces and partition faces. Figure 4 gives an overview of these TRSs.

Element TRS. The TRS of the elements consists of all the element geometric entities in the mesh. The position of a certain element in the mesh is defined by the coordinates of its nodes. As a consequence, a connectivity table that links elements to their nodes is necessary. The approximate solution polynomial on an element is determined by the states in its solution points. Consequently, a connectivity table linking the elements to states is present. Finally, a connectivity table stores indices of the neighboring faces of each element. This is necessary for the evaluation of the state and fluxes in the flux points on the faces. These connectivity tables together with the element geometric entities form the element TRS.

Inner and Boundary Face TRS. After the creation of the element TRS, the inner and boundary face TRSs are built. The position of a face in the mesh is defined by its corner nodes. As a consequence, a face-node connectivity table is made. Using the face-element connectivity table, the number of neighboring elements of a face is determined. If a

Figure 4: Overview of the TRSs that are created by the FR data structure

face has only one neighboring element, it is added to the boundary faces TRS. If it has two neighboring elements, it is added to the inner faces TRS. Since a number of different boundary conditions can apply to the boundary faces, the boundary face TRS is split up in at least the same amount of Topological Regions (TR). In this manner, a boundary condition can be applied to a subset of the boundary faces.

Finally, an additional TRS of the partition faces must be created if the solver runs in parallel. Partition faces are inner faces of the global mesh, but make up the edge of the domain of a certain CPU. Since these faces have only one neighboring element, but have no boundary condition attached to them, they must be treated separately by the FR algorithm.

4.1.2. Geometrical Data of the Elements

The FR algorithm needs to access the following geometrical data of the elements:

- type of the element: shape, geometric order, solution order and dimension;
- coordinates of the solution and flux points in the standard domain;
- local numbering of faces, flux points and solution points;
- approximate solution polynomial and its derivative.

For this purpose, a generic C++ class is defined for each type of element that provides the above information to the solver. When the data structure of the test case is built, an instance of the class is created. Since all coordinates in this class are defined in the reference domain Ω_S , as defined in section 2, it is sufficient to create only one instance of an element data class for each type of element. When the global coordinates are needed, the transformation between the reference and global domain is performed.

The distribution of the solution and flux points is a parameter of the FR method. As a consequence, the point distributions are decoupled from the generic element data classes. Point distributions are defined by the 1D coordinates of the points in the domain $[-1, 1]$. Based on these 1D coordinates, the element data classes construct the 2D or 3D coordinates of the flux and solution points depending on the type of element.

The approximate solution is only stored in the solution points. This means that in order to evaluate the approximate solution polynomial in the flux points, a Lagrange interpolation is used. Since the Lagrange base functions belonging to the i -th solution point ℓ_i are defined in the standard domain, it is sufficient to compute them once for each element type and subsequently use the polynomials for all elements of this type. In addition, the derivative of ℓ_i is computed in order to be able to determine the derivative of the discontinuous fluxes as follows:

$$\frac{\partial \hat{\mathbf{f}}^{\delta D}}{\partial \xi_j}(\boldsymbol{\xi}) = \sum_{i=1}^{N_s} \hat{\mathbf{f}}_i^{\delta D} \frac{\partial \ell_i}{\partial \xi_j}(\boldsymbol{\xi}). \quad (42)$$

The present solver can use meshes that have curved elements up to second order (Q2). For the mapping between these mesh elements and the standard domain, the approximate cross-product form as described in [65] is used. It is known that this method introduces inaccuracies when the order of the geometric mapping is larger than the solution order, for example for P1Q2. However this error tends to zero under mesh refinement. Furthermore, for all other cases (P2Q2 and higher orders of accuracy), this problem does not appear.

4.2. Implementation of the FR Algorithm

The compressible Navier-Stokes and Euler equations can be written as a system of PDEs in the conservative form as presented in equation 5. The goal of the FR solver is to compute the divergence of the reconstructed flux in each element at every time step. For this purpose, the FR algorithm is divided into different parts as shown in figure 5. As mentioned in section 1, the flux is split up in a convective and diffusive part:

$$\frac{\partial \mathbf{u}}{\partial t} = -\nabla \cdot \mathbf{f}_C(\mathbf{u}) + \nabla \cdot \mathbf{f}_D(\mathbf{u}, \nabla \mathbf{u}). \quad (43)$$

The evaluation of the physical fluxes based on $\mathbf{u}$ and $\nabla \mathbf{u}$ is performed in the physics module of the COOLFluid framework, as shown in figure 3. The convective and diffusive flux can be computed independently. As a consequence, the FR algorithm consists of two major parts: the convective algorithm and the diffusive algorithm. These subdivisions both interact with classes that define the parameters of the FR method that are specified by the user. These parameters are: the flux and solution point distribution, the interface flux scheme and the correction function type. However, the first of these parameters is used in the FR data structure as explained in the previous section. Additionally the boundary conditions as well as the shock capturing method are implemented in the FR algorithm.

Figure 5: Overview of the structure of the FR algorithm

4.3. Computation of the Convective Flux Divergence

The reconstructed convective flux in the reference domain consists of a discontinuous flux $\hat{\mathbf{f}}_C^{\delta D}$ and a correction flux $\hat{\mathbf{f}}_C^{\delta C}$ which in general are expressed as follows:

$$\hat{\mathbf{f}}_C^{\delta} = \hat{\mathbf{f}}_C^{\delta D}(\hat{\mathbf{u}}) + \hat{\mathbf{f}}_C^{\delta C}(\hat{\mathbf{u}}) = \hat{\mathbf{f}}_C^{\delta D}(\hat{\mathbf{u}}) + \sum_{i=1}^{N_f} \mathbf{h}_i(\xi) \left((\hat{\mathbf{f}}_{C,i}^{\delta I} - \hat{\mathbf{f}}_C^{\delta D}(\hat{\mathbf{u}}(\xi_i))) \cdot \hat{\mathbf{n}}_i \right), \quad (44)$$

where $\mathbf{h}_i$ is the correction field associated with the i -th flux point, N_f the total number of flux points and $\hat{\mathbf{n}}_i$ the unit normal vector to the face in flux point i . The divergence of this expression is:

$$\nabla_{\xi} \cdot \hat{\mathbf{f}}_C^{\delta} = \nabla_{\xi} \cdot \hat{\mathbf{f}}_C^{\delta D}(\hat{\mathbf{u}}) + \sum_{i=1}^{N_f} \nabla_{\xi} \cdot \mathbf{h}_i(\xi) \left((\hat{\mathbf{f}}_{C,i}^{\delta I} - \hat{\mathbf{f}}_C^{\delta D}(\hat{\mathbf{u}}(\xi_i))) \cdot \hat{\mathbf{n}}_i \right). \quad (45)$$

The computation of the discontinuous and correction flux divergence can be done separately. For the divergence of the discontinuous flux, the states in the solution points must be accessed, as shown by its dependence on $\hat{\mathbf{u}}$. Since the solution points are part of the element geometrical entities, this term is computed by performing a loop over the elements.

For the computation of the correction flux divergence, the states in the flux points ξ_i for both the left and right neighboring elements are needed, as can be seen in the second term in the RHS of equation 45: this contains the interface flux $\hat{\mathbf{f}}_{C,i}^{\delta I}$ which depends on the left and right states on a face. Since the flux points are part of the face geometric entities, this term is computed by looping over the faces. When looping over the faces, each flux point state only needs to be computed once, since each flux point is only part of one iteration of the loop. However, if a loop over the elements would be performed instead of over the faces, each face would be considered in two iterations as it belongs to two neighboring elements. This means that each flux point state would be computed twice which is inefficient.

For the computation of $\hat{\mathbf{f}}_C^{\delta D}(\hat{\mathbf{u}}(\xi_i))$ in the correction flux divergence however, the discontinuous fluxes in the solution points of an element need to be accessed. This is necessary for the extrapolation of the discontinuous fluxes to the flux points. As such, it is more efficient to compute this term in the loop over elements together with the discontinuous flux divergence. In this manner, the discontinuous fluxes in the solution points only need to be evaluated once for each element. Equation 46 summarizes which terms are computed in which loops in each solution point:

$$\text{loop over elements: } \nabla_{\xi} \cdot \hat{\mathbf{f}}_C^{\delta D}(\hat{\mathbf{u}}) - \sum_{i=1}^{N_f} \nabla_{\xi} \cdot \mathbf{h}_i(\xi) \left(\hat{\mathbf{f}}_C^{\delta D}(\hat{\mathbf{u}}(\xi_i)) \cdot \hat{\mathbf{n}}_i \right), \quad (46a)$$

$$\text{loop over faces: } \sum_{i=1}^{N_f} \nabla_{\xi} \cdot \mathbf{h}_i(\xi) \left(\hat{\mathbf{f}}_{C,i}^{\delta I}(\hat{\mathbf{u}}(\xi_i)) \cdot \hat{\mathbf{n}}_i \right). \quad (46b)$$

Note that the divergence of the correction field $\nabla_{\xi} \cdot \mathbf{h}(\xi_i)$ only has to be computed once, since it is defined within the reference domain.

4.3.1. Loop over Elements

The first term in equation 46a is the discontinuous flux divergence, which can be computed as follows in the i -th solution point:

$$\nabla_{\xi} \cdot \hat{\mathbf{f}}_C^{\delta D}(\hat{\mathbf{u}}_i) = \sum_{j=1}^{N_{dim}} \frac{\partial \hat{\mathbf{f}}_C^{\delta D}}{\partial \xi_j}(\hat{\mathbf{u}}_i) \cdot \hat{\mathbf{n}}_j = \sum_{j=1}^{N_{dim}} \sum_{k=1}^{N_s} \left(\hat{\mathbf{f}}_C^{\delta D}(\hat{\mathbf{u}}_k) \cdot \hat{\mathbf{n}}_j \right) \frac{\partial \ell_k}{\partial \xi_j}(\xi_i), \quad (47)$$

where N_{dim} is the dimensionality of the problem, N_s the number of solution points and $\hat{\mathbf{n}}_j$ the unit vector in the direction of ξ_j in the reference coordinate frame. From this equation, it is clear that the discontinuous flux divergence is computed through summation over all solution points of the projected flux multiplied by the derivative of the Lagrange base function of the corresponding solution point. The base function derivatives are defined on the reference domain and are stored in the FR data structure, as explained in section 4.1.2. The states in the solution points can easily be accessed when building an element geometric entity. The same applies to the unit axes $\hat{\mathbf{n}}_j$. As a consequence, all information required to evaluate equation 47 is available when looping over the elements.

The second term in equation 46a depends on the extrapolation of the discontinuous fluxes to the flux points. For the j -th flux point, this is equal to:

$$\hat{\mathbf{f}}_C^{\delta D}(\hat{\mathbf{u}}(\xi_j)) = \sum_{k=1}^{N_s} \hat{\mathbf{f}}_C^{\delta D}(\hat{\mathbf{u}}_k) \ell_k(\xi_j). \quad (48)$$

In this expression, every variable is now known, since the discontinuous fluxes in the solution points were evaluated to compute expression 47.

Finally, the correction field of quadrilaterals and hexahedra associated with the j -th flux point, can be expressed in terms of a scalar correction function: $\mathbf{h}_j = h_j \hat{\mathbf{n}}_j$. The scalar correction function in 2D or 3D is constructed using the tensor product of 1D correction functions, as explained in section 2. The divergence of the correction field in the i -th solution point due to the correction associated with the j -th flux point for quadrilaterals and hexahedra can simply be evaluated as follows:

$$\nabla_{\xi} \cdot \mathbf{h}_j(\xi_i) = \nabla_{\xi} \cdot (h_j(\xi_i) \hat{\mathbf{n}}_j) = \frac{\partial h_j}{\partial \xi_j}(\xi_i), \quad (49)$$

where ξ_j is the direction along the unit normal in the j -th flux point $\hat{\mathbf{n}}_j$.

In summary, the computation that is performed for the i -th solution point in the loop over elements is given in expression 50. This is the implementation of the evaluation of expression 46a:

$$\sum_{j=1}^{N_{dim}} \sum_{k=1}^{N_s} (\hat{\mathbf{f}}_C^{\delta D}(\hat{\mathbf{u}}_k) \cdot \hat{\mathbf{n}}_j) \frac{\partial \ell_k}{\partial \xi_j}(\xi_i) - \sum_{f=1}^{N_f} \frac{\partial h_f}{\partial \xi_f}(\xi_i) \sum_{n=1}^{N_s} (\hat{\mathbf{f}}_C^{\delta D}(\hat{\mathbf{u}}_n) \cdot \hat{\mathbf{n}}_f) \ell_n(\xi_f), \quad (50)$$

considering once again that ξ_f denotes the direction along the unit vector $\hat{\mathbf{n}}_f$. Note that in this expression N_f is equal to the total amount of flux points of an element. The pseudo-code of this computation is presented in algorithm 1.

4.3.2. Loop over Faces

When computing equation 46b in a loop over faces, the common interface flux needs to be evaluated. The common interface flux $\hat{\mathbf{f}}_j^{\delta I}$ can be computed if the states in the flux point of the left and right neighboring elements, respectively $\hat{\mathbf{u}}^-$ and $\hat{\mathbf{u}}^+$, are known. Therefore, the neighboring element geometric entities need to be accessed. A flux scheme such as Roe or Lax-Friedrichs flux then provides the interface flux value $\hat{\mathbf{f}}_j^{\delta I}$. The terms $\hat{\mathbf{u}}^-$ and $\hat{\mathbf{u}}^+$ are evaluated using the Lagrange interpolation, as in equation 12. Using equations 48 and 49, the final computation performed in the loop over faces for the i -th solution point of the neighboring elements becomes the following:

$$\sum_{j=1}^{N_f} \frac{\partial h_j}{\partial \xi_j}(\xi_i) ((\hat{\mathbf{f}}_{C,j}^{\delta I}(\hat{\mathbf{u}}^-(\xi_j), \hat{\mathbf{u}}^+(\xi_j)) \cdot \hat{\mathbf{n}}_j)), \quad (51)$$

where ξ_j once again denotes the direction along the unit normal vector $\hat{\mathbf{n}}_j$. Note that here N_f denotes the number of flux points on a face. Algorithm 2 presents how this is implemented in pseudo-code.

4.3.3. Treatment of the Boundary Faces

A boundary face has only one neighboring element. As a consequence, only one state (“internal” to the domain) can be evaluated in the flux points. However, there needs to be a left and right state in order to determine the interface flux. This is solved by imposing a ghost state in each flux point. Ghost states depend on the boundary condition attached to the boundary face. The way these states are computed is discussed in section 4.4.

Algorithm 2 cannot be used directly for the boundary faces. The algorithm is executed on the inner faces TRS and a separate loop over the boundary faces TRS is performed. This procedure is very similar to algorithm 2. The only differences being that the right state is now a ghost state that is computed by means of a separate boundary condition implementation and that the divergence of the correction flux is only computed for one neighbor element, i.e. the internal element attached to the boundary face.

4.3.4. Computation of the Diffusive Flux Divergence

As section 1 shows, the diffusive flux depends on both the state $\hat{\mathbf{u}}$ and the corrected gradient of the state $\hat{\mathbf{q}} = \nabla \hat{\mathbf{u}}$. Once $\hat{\mathbf{q}}$ is known, the algorithms for the computation of the convective flux divergence can be used (algorithms 1 and 2). The approximate corrected gradient in the ξ_i -direction, in the reference domain, $\hat{\mathbf{q}}_{\xi_i}^{\delta}$ is computed as follows:

$$\hat{\mathbf{q}}_{\xi_i}^{\delta} = \frac{\partial \hat{\mathbf{u}}^{\delta}}{\partial \xi_i} + \sum_{j=1}^{N_f} (\hat{\mathbf{u}}_j^{\delta I} - \hat{\mathbf{u}}^{\delta}(\xi_j)) \frac{\partial \mathbf{h}}{\partial \xi_i} \cdot \hat{\mathbf{n}}_j, \quad (52)$$

Algorithm 1 Element loop contribution to the computation of the flux divergence

```

1: procedure ELEMENTLOOP
2:   for  $iElement \leftarrow 0, N$  do
3:      $states \leftarrow$  get states of  $iElement$ 
4:     for  $iS\ state \leftarrow 0, N_s$  do
5:       for  $iDim \leftarrow 0, N_{dim}$  do
6:          $fluxes(iS\ state, iDim) \leftarrow$  get flux of  $iS\ state$  along axis  $iDim$ 
7:       end for
8:     end for
9:     for  $iFlux \leftarrow 0, N_f$  do
10:       $dim \leftarrow$  get the normal direction in  $iFlux$ 
11:      for  $iS\ state \leftarrow 0, N_s$  do
12:         $discFlux(iFlux) \leftarrow 0$  ▷ discontinuous flux
13:         $baseFct \leftarrow$  get  $\ell_{iS\ state}$  in  $iFlux$ 
14:         $discFlux(iFlux) += fluxes(iS\ state, dim) \cdot baseFct$ 
15:      end for
16:    end for
17:    for  $iS\ state \leftarrow 0, N_s$  do
18:       $result(iElement, iS\ state) \leftarrow 0$ 
19:      for  $iDim \leftarrow 0, N_{dim}$  do
20:        for  $jS\ state \leftarrow 0, N_s$  do
21:           $dBaseFct \leftarrow$  get  $\frac{\partial \ell_{jS\ state}}{\partial \xi_{iDim}}$  in  $iS\ state$ 
22:           $result(iElement, iS\ state) += fluxes(jS\ state, iDim) \cdot dBaseFct$ 
23:        end for
24:      end for
25:    end for
26:    for  $iS\ state \leftarrow 0, N_s$  do
27:      for  $iFlux \leftarrow 0, N_f$  do
28:         $dim \leftarrow$  get the normal direction in  $iFlux$ 
29:         $dh \leftarrow$  get  $\frac{\partial h_{iFlux}}{\partial \xi_{dim}}$  in  $iS\ state$ 
30:         $result(iElement, iS\ state) -= discFlux(iFlux) \cdot dh$ 
31:      end for
32:    end for
33:  end for
34:  return  $result$ 
35: end procedure

```

where $\hat{\mathbf{n}}_i$ is the unit normal vector along direction ξ_i .

The way the corrected gradient is constructed is analogous to the corrected flux. As a consequence, there is a term depending on the approximate state polynomial within an element and a term depending on the left and right state in a flux point. Following the same approach as the computation of the divergence of the convective flux, a loop over the elements and the inner and boundary faces is performed. The computation of the convective flux divergence is performed before the diffusive computation, such that the algorithm to determine the corrected gradient can be done in the loops over elements and faces of the convective algorithm. In this manner, no separate loops are needed to calculate the corrected gradients. After the execution of the convective algorithm, $\hat{\mathbf{q}}^\delta$ can be accessed to compute the diffusive flux divergence.

On the boundary, the ghost state can be used to determine the corrected gradient similarly to the approach used for the convective flux divergence. However, the method to evaluate the interface flux on the boundary must be modified. The diffusive flux depends on the state and the state gradient in the flux point. Consequently it does not

Algorithm 2 Face loop contribution to the computation of the flux divergence

```
1: procedure FACELOOP
2:   for  $iOrientation \leftarrow 0$ , maximum number of orientations do
3:     for  $iFace \leftarrow 0$ , number of faces with  $iOrientation$  do
4:        $leftNeighbor \leftarrow$  get left neighbor element of  $iFace$ 
5:        $rightNeighbor \leftarrow$  get right neighbor element of  $iFace$ 
6:       for  $iFlux \leftarrow 0, N_f$  do
7:          $leftLocalIndex(iFlux) \leftarrow$  left index of  $iFlux$  for  $iOrientation$ 
8:          $rightLocalIndex(iFlux) \leftarrow$  right index of  $iFlux$  for  $iOrientation$ 
9:          $leftState \leftarrow 0$ 
10:         $rightState \leftarrow 0$ 
11:        for  $iState \leftarrow 0, N_s$  do
12:           $iLeftState \leftarrow$  get state of left neighbor at  $iState$ 
13:           $iRightState \leftarrow$  get state of right neighbor at  $iState$ 
14:           $leftBaseFct \leftarrow \ell_{iState}$  in  $leftLocalIndex(iFlux)$ 
15:           $rightBaseFct \leftarrow \ell_{iState}$  in  $rightLocalIndex(iFlux)$ 
16:           $leftState += iLeftState \cdot leftBaseFct$ 
17:           $rightState += iRightState \cdot rightBaseFct$ 
18:        end for
19:         $interfaceFlux(iFlux) \leftarrow \hat{\mathbf{f}}_{C,iFlux}^{SI}(leftState, rightState) \cdot \hat{\mathbf{n}}_{iFlux}$ 
20:      end for
21:      for  $iState \leftarrow 0, N_s$  do
22:         $leftResult(leftNeighbor, iState) \leftarrow 0$ 
23:         $rightResult(rightNeighbor, iState) \leftarrow 0$ 
24:        for  $iFlux \leftarrow 0, N_f$  do
25:           $leftDivh \leftarrow$  get divergence of  $h_{leftLocalIndex(iFlux)}$  in  $iState$ 
26:           $rightDivh \leftarrow$  get divergence of  $h_{rightLocalIndex(iFlux)}$  in  $iState$ 
27:           $leftResult(leftNeighbor, iState) +=$ 
28:             $interfaceFlux(iFlux) \cdot leftDivh$ 
29:           $rightResult(rightNeighbor, iState) +=$ 
30:             $interfaceFlux(iFlux) \cdot rightDivh$ 
31:        end for
32:      end for
33:    end for
34:  end for
35:  return  $leftResult, rightResult$ 
36: end procedure
```

suffice to compute a ghost state, but a ghost gradient has to be computed as well. Section 4.4 describes how these are determined.

Since the correction procedure of the gradients is analogous to the correction of the fluxes, an algorithm analogous to algorithms 1 and 2 is used. The main difference being that for the correction of the fluxes the variable extrapolated to the flux points is the discontinuous flux, while for the correction of the gradients the states are extrapolated to the flux points. Now only these extrapolated states are used to compute an interface value: this interface value is in this case not dependent on the extrapolated discontinuous fluxes.

4.4. Boundary Conditions

The boundary of a mesh consists of faces that hold flux points. As such, it is precisely in these points that boundary conditions are introduced in the FR formulation. In a flux point of an inner face, a left and right state and gradient can be defined. However, there can only be one state extrapolated to a flux point that lies on a boundary face. This

is called the internal state $\mathbf{u}_{inner}$. In order to introduce the boundary condition that applies to a face, the second state and gradient in the flux point are computed based on this boundary condition. They are called ghost states $\mathbf{u}_{ghost}$ and ghost gradients $\mathbf{q}_{ghost}$.

In general, a state value in the boundary u_{bnd} can be imposed by choosing the following ghost state: (Dirichlet condition)

$$u_{ghost} = 2u_{bnd} - u_{inner}. \quad (53)$$

If a boundary condition does not impose a certain state value, the ghost state value can be set equal to the inner one, corresponding to applying a homogeneous Neumann condition for that specific value. For the gradients, the same approach is used with the exception that a boundary condition will typically impose the gradient projected on the boundary normal:

$$\mathbf{q}_{ghost} = \mathbf{q}_{inner} + 2(q_{bnd} - \hat{\mathbf{n}} \cdot \mathbf{q}_{inner}) \hat{\mathbf{n}}. \quad (54)$$

Once again, if no gradient needs to be imposed, $\mathbf{q}_{ghost}$ can simply be set to $\mathbf{q}_{inner}$.

4.5. Time Marching Procedure

The RHS of the set of equations to be solved (equation 43) is determined by discretizing the spatial domain of the problem in elements. The FR method is then used to evaluate the RHS in every solution point in each element. This discretized form of the RHS is called the residual $\mathcal{R}$:

$$\frac{\partial \mathbf{u}_{ce}}{\partial t} = \mathcal{R}(\mathbf{u}_{ce}, \mathbf{u}_{ne}), \quad (55)$$

where the subscript *ce* refers to the current element, while *ne* represents the neighboring elements. This results in a semi-discretized equation. In order to fully solve the problem, the LHS of equation 55 also needs to be discretized. This is done by dividing the time domain in discrete time steps Δt . Subsequently, a time marching scheme can determine $\mathbf{u}(t + \Delta t)$.

For implicit time stepping in addition to the residual, the Jacobian matrix of the residuals is needed. Consider as an example the backward Euler implicit time marching scheme that can be written as follows:

$$\mathbf{u}_{ce}(t + \Delta t) = \mathbf{u}_{ce}(t) + \Delta t \mathcal{R}(\mathbf{u}_{ce}(t + \Delta t), \mathbf{u}_{ne}(t + \Delta t)), \quad (56)$$

The value of Δt is determined by a *CFL* number specified by the user. At time t , the value of $\mathcal{R}(\mathbf{u}_{ce}(t + \Delta t), \mathbf{u}_{ne}(t + \Delta t))$ is not readily available since $\mathbf{u}(t + \Delta t)$ is unknown. Because $\mathcal{R}$ is non-linear for the Navier-Stokes equations, expression 56 is a non-linear equation that has to be solved during each iteration. This system of equations is linearized resulting in the following:

$$\mathbf{u}(t + \Delta t) = \mathbf{u}(t) + \Delta t \left(\mathcal{R}(\mathbf{u}(t)) + \frac{\partial \mathcal{R}}{\partial \mathbf{u}} \bigg|_t (\mathbf{u}(t + \Delta t) - \mathbf{u}(t)) \right), \quad (57)$$

where $\frac{\partial \mathcal{R}}{\partial \mathbf{u}} \bigg|_t$ is the Jacobian matrix of the residuals with respect to the state vectors, evaluated at time t . Here the subscripts *ce* and *ne* are omitted since here $\mathbf{u}$ represents the vector of all state vectors in the spatial domain. The resulting large linear system is solved with for example a Generalized Minimal Residual approach (GMRES) [66].

The residual Jacobian matrix is determined by numerically approximating the partial derivatives with a simple first-order finite difference scheme in the i -th element:

$$\frac{\partial \mathcal{R}}{\partial \mathbf{u}_i} = \frac{\mathcal{R}(\mathbf{u}_i + \varepsilon_u) - \mathcal{R}(\mathbf{u}_i)}{\varepsilon_u}. \quad (58)$$

The pseudo-code of this procedure is presented in algorithm 3. Since the contribution of all states that have an influence on the residual in a certain solution point, are taken into account, the Jacobian matrix contains both diagonal and off-diagonal blocks. This is because a perturbed state in a neighboring element will also perturb the residual in the current element and vice versa. This approach can potentially impose an important memory requirement when using large hexahedral meshes, in which case an approach where only the diagonal blocks are considered may be necessary. A concise assessment of the memory requirements of the test cases presented in the present work, is presented in section 6.

Algorithm 3 Computation of the numerical residual Jacobian

```

1: procedure COMPUTERESIDUALJACOBIAN(iElement, residual)
2:   states  $\leftarrow$  states of iElement
3:   for iState  $\leftarrow$  0,  $N_{states}$  do
4:     for iVar  $\leftarrow$  0,  $N_{var}$  do ▷ state variables
5:        $\varepsilon_u \leftarrow$  perturbation for iVar
6:       perturbedState  $\leftarrow$  states(iState)
7:       perturbedState(iVar)  $+= \varepsilon_u$ 
8:       newRes  $\leftarrow$  recompute residual with perturbedState
9:       jacobian(iState, iVar)  $\leftarrow$  (newRes - residual) /  $\varepsilon_u$ 
10:    end for
11:  end for
12:  return jacobian
13: end procedure

```

The numerical approach of computing the residual Jacobian matrix is applied to the proposed LLAV scheme. In expression 58, the value of $\mathcal{R}(\mathbf{u}_i)$ is known, since this is simply what is computed by the FR method. For the LLAV scheme, this is the artificial flux divergence. However $\mathcal{R}(\mathbf{u}_i + \varepsilon_u)$ also needs to be computed. This can be written as:

$$\mathcal{R}(\mathbf{u}_i + \varepsilon_u) = -\nabla \cdot \mathbf{f}_C(\mathbf{u}_i + \varepsilon_u) + \nabla \cdot \mathbf{f}_D(\mathbf{u}_i + \varepsilon_u, \mathbf{q}(\mathbf{u}_i + \varepsilon_u)) + \nabla \cdot \mathbf{f}_{AV}(\mathbf{q}(\mathbf{u}_i + \varepsilon_u)). \quad (59)$$

For LLAV, this results in:

$$\nabla \cdot \mathbf{f}_{AV}(\mathbf{q}(\mathbf{u}_i + \varepsilon_u)) = \varepsilon(\mathbf{u}_i + \varepsilon_u) \nabla \cdot \mathbf{q}(\mathbf{u}_i + \varepsilon_u). \quad (60)$$

$\varepsilon(\mathbf{u}_i + \varepsilon_u)$ depends on ε_0 and S , which both depend on the perturbed state. S can simply be recomputed using expression 22. For ε_0 , the characteristic flow speed is re-evaluated using the perturbed state.

In order to compute the perturbed gradient $\mathbf{q}(\mathbf{u}_i + \varepsilon_u)$, it would be possible to reapply the complete FR correction procedure to the perturbed state. However it is far more efficient to evaluate the analytical expression of the perturbed gradient, which is obtained as follows in 1D (the extension to 2D and 3D tensor elements is once again done by using the tensor product basis):

$$\hat{q}_\varepsilon^\delta = \frac{\partial \hat{u}_\varepsilon^\delta}{\partial \xi} + (\hat{u}_{L,\varepsilon}^{\delta I} - \hat{u}_{L,\varepsilon}^\delta) \frac{dg_L}{d\xi} + (\hat{u}_{R,\varepsilon}^{\delta I} - \hat{u}_{R,\varepsilon}^\delta) \frac{dg_R}{d\xi}, \quad (61)$$

where the subscript ε denotes a perturbed value. Using the Lagrange basis and the definition of the interface state (simply an average as described in section 2), this becomes:

$$\hat{q}_\varepsilon^\delta = \sum_{j=1, j \neq i}^{N_s} \hat{u}_j^\delta \frac{\partial \ell_j}{\partial \xi}(\xi) + (\hat{u}_i^\delta + \varepsilon_u) \frac{\partial \ell_i}{\partial \xi}(\xi) + \left(\frac{1}{2} \hat{u}_{ne}^\delta - \frac{1}{2} \hat{u}_{L,\varepsilon}^\delta\right) \frac{dg_L}{d\xi}(\xi) + \left(\frac{1}{2} \hat{u}_{ne}^\delta - \frac{1}{2} \hat{u}_{R,\varepsilon}^\delta\right) \frac{dg_R}{d\xi}(\xi), \quad (62)$$

where u_{ne} denotes the flux point state from the neighboring element and here the subscript i refers to the current perturbed solution point. Using the Lagrange extrapolation which computes the flux point states and the definition of the unperturbed interface state, this is written as:

$$\hat{q}_\varepsilon^\delta = \sum_{j=1}^{N_s} \hat{u}_j^\delta \frac{\partial \ell_j}{\partial \xi}(\xi) + \varepsilon_u \frac{\partial \ell_i}{\partial \xi}(\xi) + (\hat{u}_L^{\delta I} - \hat{u}_L^\delta - \frac{1}{2} \varepsilon_u \ell_i(\xi_L)) \frac{dg_L}{d\xi}(\xi) + (\hat{u}_R^{\delta I} - \hat{u}_R^\delta - \frac{1}{2} \varepsilon_u \ell_i(\xi_R)) \frac{dg_R}{d\xi}(\xi). \quad (63)$$

Using the expression of the unperturbed gradient $\hat{q}^\delta$, this finally becomes:

$$\hat{q}_\varepsilon^\delta = \hat{q}^\delta + \varepsilon_u \frac{d\ell_i}{d\xi} - \frac{\varepsilon_u \ell_i(\xi_L)}{2} \frac{dg_L}{d\xi} - \frac{\varepsilon_u \ell_i(\xi_R)}{2} \frac{dg_R}{d\xi}. \quad (64)$$

Since the unperturbed value of the gradient is known, the perturbed gradient can simply be determined by evaluating expression 64. Once $\nabla \cdot \mathbf{f}_{AV}(\mathbf{q}(\mathbf{u}_i + \varepsilon_u))$ is calculated, the residual Jacobian matrix, including the contribution of the LLAV scheme, is known. After solving the time marching equation 57, the positivity preservation scheme is applied to all elements on the newly computed state $\mathbf{u}(t + \Delta t)$.

5. Verification of the FR Solver

In this section, the developed FR solver is investigated and verified by means of various test cases dealing with both the compressible Euler and Navier-Stokes equations. First of all, the inviscid subsonic flow through a channel with a smooth sinusoid bump is studied in 2D and 3D. As a second test case governed by the Euler equations, the 2D inviscid flow around a cylinder is studied. This test case is also extended to viscous flow governed by the Navier-Stokes equations. More specifically, a laminar flow around the cylinder with stable vortices is studied in both 2D and 3D. Finally, the ability of the FR solver to accurately capture shocks is discussed based on the hypersonic flow over a cylinder and cone. All test cases that are presented in this section are solved using an implicit backward Euler time marching procedure, unless mentioned otherwise. Gauss-Legendre solution and flux points are used.

5.1. Inviscid Subsonic Flow Through a Channel with a Sinusoid Bump in 2D

In this section, the inviscid flow at $M = 0.5$ through a channel with a smooth sinusoid bump is investigated. The physical domain is bounded between the inlet ($x = 0$) and the outlet ($x = 4$), and between the bottom ($y = 0$) and the top ($y = 1$) edges. The top is straight, while the bottom consists of a smooth bump symmetrically positioned between the in- and outlet, as shown in figure 6. The bump is defined by the following expression:

$$f(x) = 0.1 + 0.1 \cos(\pi(x - 2)), \quad \text{with } x \in [1; 3]. \quad (65)$$

A slip-wall boundary condition is imposed to both the bottom and the top edge. Subsonic inlet (imposing total inlet pressure and temperature) and subsonic outlet conditions (imposing outlet pressure) are imposed such that the resulting outflow Mach number is 0.5 at zero angle of attack as proposed by [67]. The free-stream flow values are summarized in table 1.

Table 1: Free-stream values

M_∞	ρ_∞ [-]	p_∞ [-]	$v_{x,\infty}$ [-]	$v_{y,\infty}$ [-]
0.5	1	1	$0.5\sqrt{1.4}$	0

Figure 6: Channel with a smooth sinusoid bump between $x = 1$ and $x = 3$

The physical domain is discretized by means of four different grids, consisting of quadrilateral elements. These grids are composed of 3×9 , 4×13 , 5×17 , and 6×21 nodes, respectively. The first value denotes the number of nodes on the inlet and the outlet, while the second refers to the number of nodes on the top and bottom edges of the channel. The four meshes are constructed by equally distributing the nodes along the boundaries as represented in figure 7. Both straight edged elements with linear geometric mapping (Q1), as well as curvilinear elements with quadratic geometric mapping (Q2) are used. However, large errors may pollute the solution in the vicinity of the boundary if the physical boundary does not coincide with the computational boundary. Therefore, special treatment is required for the straight edged elements in order to correctly take into account the curvature of the bump as demonstrated by Krivodonova et al. [68]. As a consequence, all the simulations are run with curvilinear quadrilateral elements in order to avoid the generation of an unphysical “boundary layer” as discussed further on in this section. In this work, the elements are denoted by $PpQq$ following the convention of Bassi and Rebay [69]. Here, p refers to the order of the polynomials used to approximate the solution, and q refers to the order of the polynomials used for the geometric mapping.

Figure 7: Discretization of the channel with a bump with quadrilateral cells

The Mach contour plot obtained with P3Q2 elements on grid 4 (6×21 nodes) is shown in figure 8. For an inviscid flow, no entropy should be produced within the flow field. In addition, the geometry of the sinusoid bump channel is completely symmetrical with respect to the vertical plane at the bump ($x = 2$). Consequently, the solution for the Euler equations of this test case should also be symmetrical. From figure 8 it is clear that the flow field is smooth and is almost exactly symmetrical with respect to the bump. Figure 9 shows the logarithm of the density residual as a function of the simulation iterations. The CFL law is the following:

$$CFL(n) = \begin{cases} 0.5, & \text{if } n \leq 5, \\ \min(0.5 \cdot 2^{n-5}, 10^4), & \text{otherwise,} \end{cases} \quad (66)$$

where n is the iteration. Using implicit time stepping methods, there is no constraint on the CFL number, which greatly speeds up convergence: by being able to use large CFL numbers, convergence is reached in 21 iterations for this test case. If an explicit time stepping method is used, the CFL limit for FR is:

$$CFL_{max} = \frac{1}{P + 1}. \quad (67)$$

For the P3Q2 test case, this means that $CFL_{max} = 0.25$.

Figure 8: Mach contours for the inviscid flow through a channel with a bump obtained with P3Q2 elements on grid 4

Figure 9: Logarithm of the residual of the density for the inviscid flow through a channel with a bump obtained with P3Q2 elements on grid 4

Since the entropy should remain constant through the entire domain, the entropy error ε_s is used as an indication for the solution accuracy as proposed in [67]. The entropy error ε_s in the i -th state is defined as:

$$\varepsilon_{s,i} = \frac{s - s_i}{s}. \quad (68)$$

The exact value of the specific entropy s can be calculated using the inlet variables. Based on the total pressure p_t and the total temperature T_t imposed by the subsonic inlet boundary condition, the specific entropy s of the flow field is calculated as $s = p\rho^{-\gamma}$ where ρ and p are the inlet density and pressure. The L_1 - and the L_2 -norm of the entropy error ε_s are defined as follows:

$$L_1(\varepsilon_s) = \sum_{n=1}^N \int_{\Omega_n} |\varepsilon_s| d\Omega_n \approx \sum_{n=1}^N \sum_{j=1}^{N_q} |\hat{\varepsilon}_s| J_n \omega_j \quad \text{and} \quad L_2(\varepsilon_s) = \sqrt{\sum_{n=1}^N \int_{\Omega_n} \varepsilon_s^2 d\Omega_n} \approx \sqrt{\sum_{n=1}^N \sum_{j=1}^{N_q} \hat{\varepsilon}_s^2 J_n \omega_j}, \quad (69)$$

where $\Omega_{s,n}$ is the standard domain associated with the n -th element, J_n the Jacobian determinant of the n -th element, N_q the number of quadrature points and ω_j the j -quadrature weight. As equation 69 shows, Gauss quadrature is used to approximate the integral over the element domain as in [70].

A grid convergence study is carried out by computing the L_1 - and the L_2 -norm of the entropy error ε_s on the four successively refined grids for the P1Q2, P2Q2, P3Q2, P4Q2 and P5Q2 FR schemes. The calculation is considered to be converged when the density residual $\mathcal{R} \leq 1e-6$. In all simulations, Roe flux was used to compute the inviscid fluxes and the VCJH parameter c was set equal to the g_2 , as described in [41]. The order of accuracy of the FR solver is then computed as the slope of the linear least squares fit of

$$\log \Delta x \propto \log L(\varepsilon_s), \quad (70)$$

where L is either the L_1 - or the L_2 -norm. The characteristic spacing of the grid Δx is approximated as $N^{-1/2}$ where N is the total number of elements in the grid, following the procedure of Witherden [31]. The results of the grid convergence study are summarized in table 2. The obtained order of accuracy is approximately $P+1$ for most cases, as is expected. In the P2Q2 and P4Q2 cases a significantly higher order of accuracy than $P+1$ is reached.

It is clear that the accuracy is improved when either the polynomial order P to approximate the solution within an element is increased while maintaining the same element size, or when the mesh is refined for the same polynomial order P . The former method is commonly referred to as p -refinement, whereas the latter is known as h -refinement. Table 2 clearly illustrates the advantage of p -refinement over h -refinement. Both the L_1 - and the L_2 -norm of the entropy error obtained with P1Q2 elements on the finest grid are approximately reduced by a factor of 10 when using P4Q2 elements on the coarsest grid, even though the same number of degrees of freedom is used, namely 400. The same conclusion can be drawn when comparing the error obtained with P2Q2 elements on the finest grid to those obtained with P4Q2 elements on grid 2, while the number of DoFs is 900 for both cases. The strength of p -refinement

P	#DoF	L_1 error	L_1 order	L_2 error	L_2 order
1	64	3.16e-02	-	2.84e-02	-
	144	1.33e-02	2.13	1.58e-02	1.44
	256	7.66e-03	1.50	1.03e-02	1.92
	400	4.85e-03	2.05	7.35e-03	1.51
2	144	5.19e-03	-	5.53e-02	-
	324	1.46e-03	3.13	1.84e-03	2.72
	576	5.88e-04	3.16	7.60e-04	3.06
	900	2.99e-04	3.04	3.83e-04	3.07
3	256	1.05e-03	-	1.43e-03	-
	576	2.00e-04	4.08	2.85e-04	3.98
	1024	6.98e-05	3.67	1.05e-04	3.49
	1600	3.31e-05	3.34	5.12e-05	3.20
4	400	3.26e-04	-	6.6e-04	-
	900	3.22e-05	5.70	7.16e-05	5.42
	1600	6.11e-06	5.78	1.65e-05	5.11
	2500	2.35e-06	4.28	6.15e-06	4.42
5	576	2.38e-04	-	5.96e-04	-
	1296	2.15e-05	5.93	6.05e-05	5.64
	2304	3.35e-06	6.46	1.07e-05	6.02
	3600	1.06e-06	5.04	3.36e-06	5.19

Table 2: Results of the grid convergence study for the 2D inviscid flow through a channel with a bump (Q2 curvilinear elements) using the L_1 - and L_2 -norm of the entropy error ε_s as an indication for the solution accuracy

is clearly illustrated in figure 10, which shows the Mach contours obtained with second- (P1Q2), fourth- (P3Q2), and sixth-order (P5Q2) FR schemes on the coarsest grid.

Figure 10: Mach contours for the inviscid flow through a channel with a smooth bump on the coarsest grid (3×9 nodes) under p -refinement

5.1.1. Treatment of the Boundary

For high-order methods, the mesh approximation of the physical boundary has a large influence on the solution as discussed in [68, 71, 72]. If the order of the geometric approximation of the boundary is too low in comparison to the solution polynomial order, the solution accuracy drops significantly and unphysical phenomena for Euler problems, such as spurious wakes and boundary layers, can form. To illustrate this, consider the P3Q2 and P3Q1 solutions on grid 4 using Lax-Friedrichs flux. For the solution where the mesh has a geometric order of Q1, it is clear that an unphysical “entropy layer” is generated behind the bump as shown in figure 11. This layer however is not present in the Q2 solution. As a consequence, the accuracy obtained with curvilinear Q2 elements is significantly higher than the accuracy obtained with straight edged Q1 elements:

$$L_1(\varepsilon_s(t_\infty))_{Q1} = 2.88\text{e-}03, \quad L_2(\varepsilon_s(t_\infty))_{Q1} = 6.48\text{e-}03, \quad (71)$$

$$L_1(\varepsilon_s(t_\infty))_{Q2} = 1.33\text{e-}05, \quad L_2(\varepsilon_s(t_\infty))_{Q2} = 2.12\text{e-}05. \quad (72)$$

Therefore, the use of straight edged Q1 elements requires special treatment in order to take into account the curvature of the bump. A possible solution for this is the introduction of analytical boundary normals, as applied in [68, 72]. It consists of using the normal to the exact analytical form of the physical boundary for the flux projection on the boundary instead of the constant normal along the straight edged element. This is illustrated in figure 12. This approach gives a similar result as the use of curvilinear elements.

Figure 11: Mach contours obtained with P3Q1 and P3Q2 elements for grid 4

Figure 12: Schematic presentation of the use of analytical boundary normals

5.1.2. Parameter Study

This section investigates the influence of the parameters of the VCJH c -parameter as well as the flux and solution point distribution on the accuracy of the obtained solution for the 2D channel with a sinusoid bump test case. In order to investigate the influence of a parameter, a reference test case is chosen. This case is the 2D channel with a sinusoid bump with geometric order Q2 on grid 3. Unless mentioned otherwise, the FR parameters are: Roe flux, the c_{g_2} value of the c parameter and Gauss-Legendre points for the solution and flux point distribution.

The VCJH correction field for quadrilaterals and hexahedra is fully defined by a single parameter c . The FR method with a VCJH correction function is proven to be energy stable for $c_- < c < \infty$ [73]. Let c_{min} be the minimum value for which the developed FR solver does not diverge for the flow through the 2D channel with a sinusoid bump. Consequently, the present FR solver converges for this test case if $c_{min} \leq c < \infty$. This value is determined experimentally by progressively lowering c until no convergence is reached. Table 3 shows the values of c_- and c_{min} for orders P1 through P5. As expected, the values of c_{min} are close to those of c_- .

P	c_-	c_{min}
1	-6.667e-1	-6.402e-1
2	-4.444e-2	-4.345e-2
3	-1.270e-3	-1.232e-3
4	-2.016e-5	-1.926e-5
5	-2.036e-7	-1.912e-7

Table 3: Values of c_- and c_{min} for orders P1 through P5

Once c attains a certain value c_{large} , the form of the correction function does not significantly change anymore. When c is further increased beyond c_{large} the L_1 -norm of ε_s reaches a constant value and the accuracy stagnates as shown in figure 13. Figure 14 shows the variation of the L_1 -norm of the entropy error ε_s for different values of c for orders P1, P3 and P5. It is clear that the L_1 -norm of the entropy error ε_s reaches a minimum for a value of c close to c_{DG} , which is in line with the results in [41].

Figure 13: Stagnation of the L_1 -norm of the entropy error ε_s as c approaches c_{large}

Figure 14: Influence of the VCJH parameter c on $L_1(\varepsilon_s)$

As proposed by Jameson, Vincent and Castonguay in [41, 74], the coordinates of the solution and flux points have a large influence on non-linear aliasing driven instabilities. The use of quadrature points such as Gauss-Legendre is advised. In the present FR solver, the Gauss-Legendre points, Gauss-Lobatto points and an equidistant set of points have been implemented. The reference test case is executed with these three types of point distributions. The solution and flux point distributions are always chosen to be the same. The L_1 -norm of the entropy error for these distributions is shown in table 4.

P	Equidistant	Gauss-Legendre	Gauss-Lobatto
1	5.94e-03	7.66e-03	1.48e-02
2	3.21e-04	5.88e-04	1.17e-03
3	6.27e-05	6.98e-05	1.11e-04
4	7.88e-06	6.11e-06	1.25e-05
5	-	3.35e-06	7.14e-06

Table 4: L_1 norm of the entropy error for different convective interface flux schemes for orders P1 through P5

Although the performance of the equidistant set is similar to Gauss-Legendre for lower orders, for higher orders it is considerably worse. For order P5, the FR method with an equidistant set of solution and flux points was not able to converge. This is consistent with the observation that quadrature points should be used to minimize aliasing driven instabilities [41, 74]. Furthermore, it can be noted that the Gauss-Legendre distribution is an order of magnitude more accurate than the Gauss-Lobatto points.

5.2. Inviscid Subsonic Flow Through a Channel with a Sinusoid Bump in 3D

In order to verify the FR solver in 3D, the 2D grids presented in the previous section are extruded to 3D while keeping the geometry of the xy -plane constant. The length of the domain in the z -direction equals one. The 3D meshes consist of $3 \times 9 \times 3$, $4 \times 13 \times 4$, $5 \times 17 \times 5$, and $6 \times 21 \times 6$ nodes where the last number refers to the number of nodes in the z -direction. The grids consist of curvilinear Q2 hexahedral elements.

The boundary conditions used in this test case are similar to the 2D case: an inlet and outlet with the same parameters are used, however, the inlet and outlet are planes in the 3D case instead of edges. For the bottom and top wall, a slip-wall boundary condition is used. The same applies to the newly created walls by extruding the 2D mesh. The same FR parameters and CFL law as for the 2D case were used. Figure 15 presents the Mach contours on grid 4 for P3Q2 elements. As expected, this solution is equivalent to the 2D Mach contours shown in figure 8 in the xy -plane, while the solution is constant in the z -direction.

A convergence study is performed using the same approach as presented in section 5.1. For the 3D case, the characteristic grid spacing is approximated by $\log \Delta x = \log(N^{-1/3})$. Table 5 presents the results of this convergence study. As expected, the results are similar to the 2D convergence study, which are shown in table 2.

Figure 15: Mach contours obtained with P3Q2 elements on grid 4 in 3D

5.3. Inviscid Subsonic Flow Around a Cylinder in 2D

In this section, the inviscid flow at $M_\infty = 0.38$ around a circular cylinder with a diameter equal to one is studied. The obtained results are compared to reference solutions provided by Bassi and Rebay [69], Krivodonova and Berger [68], and Van den Abeele [72]. The freestream values are given in table 6.

Table 6: Free-stream values

M_∞	ρ_∞ [-]	p_∞ [-]	$v_{x,\infty}$ [-]	$v_{y,\infty}$ [-]
0.38	1	1	$0.38 \sqrt{1.4}$	0

The physical domain is discretized by means of four O-grids consisting of quadrilateral elements, as shown in figure 16. These grids are equivalent to those used by Van den Abeele [72] and consist of 16×5 , 32×9 , 64×17 , and 128×33 nodes, respectively. The first value denotes the number of nodes on the inner and outer circle, while the second number refers to the number of concentric circles. The nodes in the circumferential direction are equally

P	#DoF	L_1 error	L_1 order	L_2 error	L_2 order
1	256	3.16e-02	-	2.84e-02	-
	864	1.33e-02	2.13	1.58e-02	1.44
	2048	7.49e-03	2.00	9.87e-03	1.64
	4000	4.72e-03	2.07	7.11e-03	1.47
2	864	5.19e-03	-	5.54e-03	-
	2916	1.46e-03	3.13	1.84e-03	2.72
	6912	4.84e-04	3.84	6.60e-04	3.56
	13500	2.53e-04	2.91	3.35e-04	3.04
3	2048	1.06e-03	-	1.45e-03	-
	6912	2.02e-04	4.10	2.88e-04	3.99
	16384	7.05e-05	3.65	1.03e-04	3.57
	32000	3.31e-05	3.38	5.12e-05	3.14
4	4000	5.34e-04	-	8.38e-04	-
	13500	6.11e-05	5.35	1.04e-04	5.14
	32000	8.35e-06	6.92	2.03e-05	5.68
	62500	2.51e-06	5.38	6.13e-06	5.38

Table 5: Results of the grid convergence study for the 3D inviscid flow through a channel with a bump using the L_1 - and L_2 -norm of the entropy error ε_s as an indication for the solution accuracy

distributed over the cylinder and the far field boundary, while those in the radial direction are distributed by means of a geometric progression using the Gmsh mesh generator [4]. In this manner, the radii of the circles in the finest grid are given by the following expression:

$$r_i = r_0 \left(1 + \frac{2\pi}{128} \sum_{n=0}^{i-1} \zeta^n \right), \quad i = 0, \dots, 32, \quad (73)$$

where $r_0 = 0.5$ is the radius of the cylinder. The parameter $\zeta = 1.1648336$ is chosen such that the radius of the outer circle r_{32} , associated with the far field boundary, is 20. This grid is then successively coarsened in order to obtain the three other grids, as shown in figure 16. A slip-wall boundary condition is introduced at the cylinder wall ($r_{\text{cylinder}} = 0.5$). At the boundary $r_{\text{farfield}} = 20$, a far field boundary condition is imposed.

As demonstrated by Bassi and Rebay [69], curvilinear Q2 elements with quadratic geometric mapping are mandatory in order to take into the account the curvature of the cylinder. In the case of straight edged elements, the solution is heavily polluted in the vicinity of the cylinder wall due to the fact that the computational boundary substantially differs from the physical boundary. As a result, an unphysical “wake” is generated at the downstream side of the cylinder leading to inaccurate solutions. As outlined in the previous section, special measures need to be taken in order to avoid this spurious entropy production.

The Mach isolines obtained with P1Q2 elements are plotted in figure 17 for the four successively refined grids. Figure 18 shows the Mach isolines resulting from second- (P1Q2), fourth- (P3Q2) and sixth-order (P5Q2) FR schemes on the coarsest grid. The VCJH parameter c is chosen such that the SD method is retrieved, as used by Van den Abeele in [72]. The obtained solutions show great similarities to those in [68, 72, 69]. The same CFL law as for the sinusoid bump test case is used. By comparing figure 17 to 18, it is once again clear that p -refinement is more favorable than h -refinement. The Mach contours obtained with P5Q2 elements on the coarsest grid, which requires 1600 DoFs, compare well with those obtained with P1Q2 elements on the finest grid, with 16384 DoFs.

Figure 16: Meshes used for the inviscid flow around a cylinder test case, consisting of Q2 quadrilaterals, same meshes as used in [72]

Figure 17: Mach contours for the inviscid flow around a cylinder obtained with P1Q2 elements under h -refinement

Figure 18: Mach contours for the inviscid flow around a cylinder obtained on the coarsest grid (16×5 nodes) under p -refinement

5.4. Viscous Subsonic Flow Around a Cylinder in 2D

The viscous flow around cylinders is greatly influenced by the Reynolds number:

$$\text{Re} = \frac{\rho v_\infty D}{\mu}, \quad (74)$$

where v_∞ is the flow speed far away from the cylinder, D the diameter of the cylinder and μ the dynamic viscosity. Based on the value of Re , different flow regimes can be defined. The flow is characterized as laminar up to $\text{Re} = 10^3$. In the vicinity of this Re -number, there is a transitional phase to a fully turbulent flow. For a very low Reynolds number, e.g. $\text{Re} < 1$, the flow is symmetrical and no separation occurs.

For $1 < \text{Re} < 40$ a stable re-circulation bubble is formed behind the cylinder. In this manner, two steady vortices appear at its trailing edge. For $40 < \text{Re} < 10^3$ vortices are shed behind the cylinder and a von Karman street is formed. [75, 76, 77]

In order to compute a steady state solution that contains stable vortices, a Reynolds number of 40 is chosen. The free-stream Mach number equals $M_\infty = 0.15$. The Prandtl number defined as $\text{Pr} = \frac{\mu c_p}{\kappa}$ where κ is the Fourier heat conduction coefficient and c_p the specific heat capacity. For this test case Pr is taken equal to $\text{Pr} = 0.72$, the standard value for air.

The cylinder has a diameter of $D = 1$ and is considered as an adiabatic no-slip wall. The computational domain is equal to $[-20, 20] \times [-20, 20]$. The outer circle that limits the computational domain is associated with a far field boundary condition. The free-stream parameters that are imposed at this boundary are given in table 7. The present test case is similar to the test case used to verify the SD solver in [72] and the FR solver in [36]. The test case is run on grid 3 that was used for inviscid flow over a cylinder, shown in figure 16c.

Table 7: Free-stream values

M_∞	ρ_∞ [-]	p_∞ [-]	$v_{x,\infty}$ [-]	$v_{y,\infty}$ [-]
0.15	1	1	$0.15\sqrt{1.4}$	0

The Mach contours of the solution for order P1, P3 and P5 are shown in figure 19. Streamtraces are plotted as well to visualize the re-circulation zones. For each order, the same origin points of the streamtraces are used. It is clear that all orders predict the stable vortices behind the cylinder. It is also clear that towards higher orders, the quality of the solution increases: the resolution of the vortices is more precise. For order P1, a single streamtrace inside the vortex does not describe a clearly defined circular path. However, for increasing orders, the path becomes better defined. Note that these solutions are consistent to the ones obtained in [36] and [72]. Figure 20 shows the logarithm of the density residual as a function of the simulation iterations. The CFL law is the following:

$$CFL(n) = \begin{cases} 0.5, & \text{if } n \leq 10, \\ \min(0.5 \cdot 2^{n-10}, 10^4), & \text{otherwise.} \end{cases} \quad (75)$$

where n is the iteration. Using this CFL law, the present test case converges in 26 iterations. Once again if an explicit time stepping method is used, the CFL limit for FR would be $CFL_{\max} = 0.25$ in this P3Q2 case with an even more severe limitation for higher orders.

Figure 19: Mach contours for the viscous flow over a cylinder for the grid in figure 16c

Figure 20: Logarithm of the residual of the density as a function of the iteration for the P3Q2 viscous 2D cylinder test case

5.4.1. Viscous Subsonic Flow Around a Cylinder in 3D

Similarly to the approach used for the verification of inviscid 3D flow, an extrusion is made of the coarsest cylinder mesh shown in figure 16a. In the newly formed z -direction, two elements are added of length one. The total dimension in the z -direction thus equals two. This results in a 3D mesh with $16 \times 5 \times 3$ nodes in the x -, y -, and z -direction respectively. The cylinder is still considered as an adiabatic no-slip wall. Similar to the 2D case, a far field boundary condition is imposed to the outer border of the computational domain. The newly created top and bottom faces that are parallel to the xy -plane are associated with a mirror boundary condition. This is because the flow should be symmetrical with respect to the xy -plane at these boundaries. The same CFL law as the 2D case is used.

For 3D viscous flow around a cylinder, additional instabilities can occur that induce a transition from a 2D to a 3D flow. In [78], two modes of transition are proposed: one for $Re \approx 190$ and one for $Re \approx 260$ which are both higher than the considered Reynolds number. Figure 21 shows the Mach contours and streamtraces for order P2 and P4. Indeed, it is clear that the flow is constant in the z -direction, as the streamtraces are confined in a vertical plane. In addition, the streamtraces are similar to the ones obtained in the 2D test case. The steady vortices behind the cylinder are thus accurately predicted in the 3D case as well. Although the same mesh is used for the P2 and P4 solution, it is clear that the latter results in a better resolution of the steady vortices, which is also consistent with the 2D case.

Figure 21: Mach contours for the viscous flow over a cylinder in 3D for orders P2Q2 and P4Q2 for the grid in figure 16a extruded in the z -direction

5.5. Viscous Hypersonic Flow over a Cylinder in 2D

The proposed shock capturing scheme was tested for a hypersonic cylinder test case, presented in [79]. The considered geometry consists of a cylinder with radius 1m. The free-stream values imposed are presented in table 8.

Table 8: Free-stream values

M_∞	$\rho_\infty [kg/m^3]$	$V_\infty [m/s]$	$T_\infty [K]$	$T_{wall} [K]$
17.637	0.001	5000	200	500

The mesh used is presented in figure 22. The mesh consists of curved quadrilateral elements of second-order and contains 4000 elements. The inlet boundary is imposed as a Dirichlet boundary condition that sets the states to the free-stream values. For the outlet, a supersonic outlet boundary condition is used. The cylinder is represented by an iso-thermal no-slip wall.

The FR method is used with a VCJH scheme for the correction functions. The distribution of both the solution and flux points is the Gauss-Legendre distribution. For the convective interface fluxes, the AUSM⁺ scheme is used. For the LLAV scheme, κ is set to 4 and c to 0.1.

Figure 22: Mesh used for the hypersonic cylinder test case

The pressure contours and artificial viscosity field for the P2Q2 solution are presented in figure 23. It is clear that the LLAV method is only active in the shock area. Figure 24 shows the wall pressure of the P1Q2 and P2Q2 solution compared to the pressure found in [79]. For P1 the pressure is overestimated, while for P2 the pressure coincides with the reference.

Figure 23: Flow field of the hypersonic cylinder test case for third-order of accuracy

Figure 25 shows the influence of the LLAV parameter c on the computed shock thickness. This result is used to calibrate the relation between c and δ_s^* . From equation 34:

$$\alpha = \frac{c M_\infty}{\delta_s^* (M_\infty - 1)} \quad (76)$$

For $c = 0.2$, the shock is spread over about 8 elements, as shown in figure 25c. Substituting $c = 0.2$ and $\delta_s^* = 8$ in equation 76, gives $\alpha = 0.027$. Figure 25b shows that for $c = 0.1$, the shock is spread over about 4 elements, which gives the same result for α . For $c = 0.05$ the shock is spread over about two elements, as shown in figure 25a, which again confirms $\alpha = 0.027$. However, for this last value of c , the positivity preservation scheme remains active until convergence.

Figure 24: Wall pressure for the hypersonic cylinder test case

Figure 25: Shock thickness for different values of c , density ρ is shown

5.6. Viscous Hypersonic Flow over a Cone in 2D

The second test case considered to test the LLAV shock capturing scheme is the flow of nitrogen gas over a rounded cone. The considered geometry consists of a 25° half-angle cone with a rounded nose of radius 1mm. The total length of the cone in the x -direction considered is 9cm. The free-stream values imposed are presented in table 9. This corresponds to Run 35 conditions as presented in [80].

Table 9: Free-stream values

M_∞	ρ_∞ [kg/m ³]	V_∞ [m/s]	T_∞ [K]	T_{wall} [K]
11.3	0.5515	2713	138.9	296

The mesh used is presented in figure 26. The mesh consists of curved quadrilateral elements of second-order and contains 2000 elements. The inlet boundary is imposed as a Dirichlet boundary condition that sets the states to the free-stream values. For the outlet, a supersonic outlet boundary condition is used. The cone is represented by an iso-thermal no-slip wall.

The FR method is used with a VCJH scheme for the correction functions. The distribution of both the solution and flux points is the Gauss-Legendre distribution. For the convective interface fluxes, the AUSM⁺ scheme is used. For the LLAV scheme, κ is set to 2 and δ_s^* to 2. In order to simulate the flow over a cone on a 2D mesh, axisymmetric source terms are added to the Navier-Stokes equations. Appendix A describes in detail how these source terms are discretized using the FR method and how they are evaluated for the axisymmetric case.

Figure 26: Mesh used for the hypersonic cone test case

The pressure contours are presented in figure 27a. Figure 27b shows the artificial viscosity distribution. It is clear once again that the LLAV method is only active in the shock area.

Figure 27: Flow field of the hypersonic cone test case for fourth-order of accuracy

Figure 28 presents the density residual history and the CFL number for both the P1 and P3 cases. For hypersonic test cases the CFL number is chosen smaller than for the previous subsonic cases in order to preserve stability. However, for both P1 and P3 the CFL number is still much larger than the CFL limit explicit time marching methods would impose: $CFL_{max} = 0.5$ and $CFL_{max} = 0.25$ for respectively P1 and P3. For the P3 case it should be noted that at iteration 660 the AV field was frozen (ε is no longer recomputed at each iteration) in order to speed up convergence.

Figure 28: Logarithm of the density residual and CFL number as a function of the iteration for the P1 and P3 cone test case

Figure 29a shows the heat flux to the surface of the cone for different heights of the first cell near the wall. This was done for P1 solutions and for the previously mentioned LLAV settings. As the mesh gets finer near the wall, the heat flux coincides with the reference heat flux. The reference was obtained as the grid-converged heat flux solution for the second-order FV solver within COOLFluid, which has been well validated for hypersonic test cases [81]. As such it is clear that if the mesh is fine enough, the FR solver computes the correct heat flux.

Figure 29b shows the influence of the LLAV κ - and c -parameters on the heat flux. It is clear that if κ is increased, the heat flux that is found decreases dramatically. This is due to the fact that for large values of κ AV is added outside of the shock area, and mainly in the boundary layer where strong gradients can be found. This causes a tendency to “smear out” the boundary layer, drastically decreasing the heat flux. The values for the c -parameter considered: (0.05;0.1;0.5), correspond to values of the characteristic shock thickness δ_s^* of respectively (2;4;20). For increasing the c -parameter, the same effect is found but much less dramatically. When c is increased, more AV is added, as ε_0 scales with c , which causes the heat flux to drop slightly. Furthermore, it is possible to retrieve almost exactly the reference heat flux in this P1 test case with a specific set of (κ, c) -values: $\kappa = 1.3$ and $c = 0.07$. In that case the overestimation of the heat flux is countered by the spurious artificial viscosity. In this way these two errors cancel each other out. Finally, figure 29c shows three cases with the same amount of degrees of freedom, namely 8000 (for P1 the mesh in figure 26 is used). These cases are a P1 FR solution, a P3 FR solution (both with the LLAV settings used for the solution in figure 27) and a second-order FV solution. Despite having the same amount of degrees of freedom, the P3 FR solution is the most accurate. This illustrates the higher computational efficiency per degree of freedom of high-order methods.

Figure 29: Heat flux for different cases

6. Performance

In order to assess the performance of the present solver, its scalability for different orders is investigated for both explicit and implicit time marching. For the Euler equations, the 2D inviscid flow over a cylinder described in section 5.3 as well as the 3D inviscid flow in a channel with a sinusoid bump described in section 5.2 are taken as benchmark cases. For the Navier-Stokes equations, the viscous flow over a cylinder test case described in sections 5.4 and 5.4.1 is considered. All simulations were run on a Dell XPS 15 notebook with an Intel Core i7 quad-core processor with a clock speed of 3.8GHz and 16GB of RAM. In this section only single-core performance is analyzed. The parallelization of the present solver is handled by the COOLFluid framework. The parallelization of COOLFluid has been analyzed extensively in [38, 58, 59], as such it will not be addressed in the present paper.

Figure 30 presents the time needed per iteration per element in the mesh for the 2D cylinder test cases. For these simulations, the grid shown in figure 16c was used. All cases were run for 10000 iterations with a forward Euler time marching scheme on a single core over which an average was taken to obtain the time per iteration. The execution time for order P1-P9 is compared with the execution time obtained when the cases are run in P1 but with the same

amount of degrees of freedom. As such, the curves labeled “P1” can be approximated as follows:

$$t_{P1}(p) = \frac{N_{DoF,Pp}}{N_{DoF,P1}} t_{P1}(1), \quad (77)$$

where $N_{DoF,Pp}$ is the amount of degrees of freedom for order p and $t_{P1}(p)$ is the execution time of P1 for the same amount of degrees of freedom as the order Pp case. By comparing the execution time of each order on the same mesh with the P1 execution time for the same amount of degrees of freedom, it can be determined whether higher-orders require a larger or smaller execution time per degree of freedom.

From figure 30, it is clear that for an explicit method, higher orders outperform P1 with the same amount of degrees of freedom in terms of execution time per degree of freedom. This is true for both Euler and Navier-Stokes (which is abbreviated to NS in the figure).

Figure 30: Execution time per iteration per element for the 2D Euler and Navier-Stokes cylinder test cases for explicit time marching

Figure 31 presents the time needed per iteration per element for the 2D test cases with implicit time marching. For these cases, only the time to assemble the linear system for the backward Euler method is considered, not the time it takes PETSc [82] to solve the system, since this is independent of the implementation of the FR algorithm. Furthermore, for the implicit cases, 1000 iterations were run of which an average was computed to obtain the time per iteration. Due to the fact that a numerical Jacobian method is used, higher orders perform worse than simply increasing

Figure 31: Execution time per iteration per element for the cylinder test cases for implicit time marching

the amount of degrees of freedom in a low order. However up to order P5, the performance of high-order compared to low-order for the same amount of degrees of freedom is similar. Only for P6 and higher, the increased cost of having

to recalculate the residual based on a perturbed state in an element with a large amount of solution points, becomes significant. This scaling problem could potentially be resolved by implementing an analytical Jacobian computation instead of the current numerical approach.

For the 3D inviscid case, the flow through a channel with a sinusoid bump is considered, using the mesh 4x13x4 presented in section 5.2. For the viscous case, the flow over a cylinder was considered with the mesh presented in section 5.4.1. Figure 32 presents the time needed per iteration per element using the forward Euler method for these two test cases. The performance is only shown up to P8 in order to avoid having to run a 3D P9 case with a very long execution time. It is clear that the same tendency as in 2D is observed: for explicit time marching, the high-order cases outperform the low-order ones with the same amount of degrees of freedom. However, for P7 and P8 in the viscous case, the execution time is slightly larger than the P1 case with the same amount of degrees of freedom.

Figure 32: Execution time per iteration per element for the 3D Euler and Navier-Stokes test cases for explicit time marching

Figure 33 presents the time needed per iteration per element for the 3D test cases with implicit time marching up to P5. Once again only the time needed to assemble the linear system of the backward Euler method was considered. Compared to the 2D case, high-order performs worse in respect to the P1 case. Already for P4 and higher the execution time is significantly larger as compared to P1 with the same amount of degrees of freedom. This could once again potentially be solved by using a analytical approach to compute the Jacobian instead of a numerical one.

Figure 33: Execution time per iteration per element for the 3D test cases for implicit time marching

Finally, table 10 shows the memory requirement per element for the different cases considered above for P1 and P5.

Table 10: Memory usage per element in kB

		Euler		Navier-Stokes	
		Explicit	Implicit	Explicit	Implicit
2D	P1	10.5	20.8	10.9	23.6
	P5	14.3	478	17.2	488
3D	P1	31.3	293	38.8	1240
	P5	257	5480	298	26100

7. Conclusion

The implementation of a novel FR solver has been presented in detail in the present paper. The solver handles 2D and 3D flows governed by the Euler or Navier-Stokes equations on unstructured meshes of quadrilaterals or hexahedra. The solver is fully implicit and parallel. In particular, a novel positivity preserving shock capturing scheme based on the LLAV approach has been developed and implemented, allowing the method to tackle high-speed flow simulations up to Mach 17.6 for the first time to the authors' knowledge. An extensive verification of the developed code has been presented for 2D, 3D, inviscid, viscous, subsonic and hypersonic test cases. As the main achievement of this work the first open-source implicit FR solver able to handle viscous hypersonic flows has been developed.

8. Acknowledgments

The research of Ray Vandenhoeck is supported by SB PhD fellowship 1S19918N of the FWO.

Appendix A. FR Method Including the Axisymmetric Source Term

This appendix describes the discretization of source terms within the FR solver as well as the implementation of the axisymmetric source term. This source term is added to the 2D Navier-Stokes equations for the hypersonic cone test case described in section 5.6.

Appendix A.1. Discretization of Source Terms

Consider the advection-diffusion equations governing the fluid motion with a source term $\mathbf{s}(\mathbf{u})$:

$$\frac{\partial \mathbf{u}}{\partial t} = -\nabla \cdot \mathbf{f}(\mathbf{u}) + \mathbf{s}(\mathbf{u}). \quad (\text{A.1})$$

Applying the space discretization as described in section 2, equation A.1 is considered on the i -th solution point of an element of the computational domain:

$$\left(\frac{\partial \mathbf{u}}{\partial t} \right)_i = -(\nabla \cdot \mathbf{f}(\mathbf{u}))_i + (\mathbf{s}(\mathbf{u}))_i. \quad (\text{A.2})$$

The approximate state vector, $\mathbf{u}^\delta$ based on a Lagrange polynomial basis, is used to approximate $\mathbf{u}$ in equation A.2:

$$\left(\frac{\partial \mathbf{u}^\delta}{\partial t} \right)_i = -(\nabla \cdot \mathbf{f}(\mathbf{u}^\delta))_i + (\mathbf{s}(\mathbf{u}^\delta))_i. \quad (\text{A.3})$$

Transforming the equation A.3 into the standard computational domain, using equation 10, the following is found:

$$\left(\frac{\partial \hat{\mathbf{u}}^\delta}{\partial t} \right)_i = -(\nabla \cdot \mathbf{f}(\hat{\mathbf{u}}^\delta))_i + \underline{(\mathbf{s}(\hat{\mathbf{u}}^\delta))}_i. \quad (\text{A.4})$$

The discretization of the source terms, i.e. the underlined term in equation A.4, is computed under the following assumption, which expresses that the source term can simply be computed based on the local value of the state:

$$(\underline{\mathbf{s}(\hat{\mathbf{u}}^\delta)})_i = \mathbf{s}(\hat{\mathbf{u}}_i^\delta). \quad (\text{A.5})$$

Consequently, the FR solver computes the source terms using the approximate state vector $\hat{\mathbf{u}}_i^\delta$ and adds them to the RHS of the equation A.1. In practise the source term is evaluated in each solution point using the approximate state polynomial in that solution point and added to the RHS in that solution point. For the Jacobian, the same numerical flux approach is used as for the convective and diffusive flux Jacobian.

Appendix A.2. Axisymmetric Source Term

The first formulation presented in [58] is used to compute axisymmetric test cases. In this formulation, the convective and diffusive fluxes remain the unchanged. However, a source term is added to the conservative form of the equations (5) that has the following form:

$$\mathbf{s} = (0, 0, p - \tau_{\theta\theta}, 0)^T, \quad (\text{A.6})$$

where $\tau_{\theta\theta}$ is the viscous stress component in the circumferential direction θ . $\tau_{\theta\theta}$ is computed using the following relation:

$$\tau_{\theta\theta} = -\frac{2}{3}\mu \left(\frac{\partial v_x}{\partial x} + \frac{\partial v_r}{\partial r} - 2\frac{v_r}{r} \right), \quad (\text{A.7})$$

where r is the radius, which is chosen equal to the y -coordinate. The partial derivatives in equation A.7 can simply be computed using the corrected gradient $\mathbf{q}^\delta$. The corrected gradient is calculated for the diffusive flux evaluation and as such is already known when evaluating the axisymmetric source term.

References

- [1] A. Sheshadri, An analysis of stability of the flux reconstruction formulation with applications to shock capturing, Ph.D. thesis, Stanford University (2016).
- [2] P. E. Vincent, A. Jameson, Facilitating the adoption of unstructured high-order methods amongst a wider community of fluid dynamicists, *Mathematical Modelling of Natural Phenomena* 6 (3) (2011) 97–140.
- [3] H. T. Huynh, Z. J. Wang, P. E. Vincent, High-order methods for computational fluid dynamics: a brief review of compact differential formulations on unstructured grids, *Computers & Fluids* 98 (2014) 209–220.
- [4] C. Geuzaine, J.-F. Remacle, Gmsh: A 3-D finite element mesh generator with built-in pre-and post-processing facilities, *International journal for numerical methods in engineering* 79 (11) (2009) 1309–1331.
- [5] S. K. Lele, Compact finite difference schemes with spectral-like resolution, *Journal of computational physics* 103 (1) (1992) 16–42.
- [6] T. Pulliam, High order accurate finite-difference methods: as seen in OVERFLOW, in: 20th AIAA Computational Fluid Dynamics Conference, 2011, p. 3851.
- [7] W. D. Henshaw, D. W. Schwendeman, Moving overlapping grids with adaptive mesh refinement for high-speed reactive and non-reactive flow, *Journal of Computational Physics* 216 (2) (2006) 744–779.
- [8] A. Lani, B. Sjögreen, H. C. Yee, W. D. Henshaw, Variable high-order multiblock overlapping grid methods for mixed steady and unsteady multiscale viscous flows, part II: hypersonic nonequilibrium flows, *Communications in Computational Physics* 13 (2) (2013) 583–602.
- [9] T. Barth, P. Frederickson, Higher order solution of the Euler equations on unstructured grids using quadratic reconstruction, in: 28th aerospace sciences meeting, 1990, p. 13.
- [10] T. J. Barth, H. Deconinck, High-order methods for computational physics, Vol. 9, Springer Science & Business Media, 2013.
- [11] M. Delanaye, Y. Liu, Quadratic reconstruction finite volume schemes on 3D arbitrary unstructured polyhedral grids, in: 14th Computational Fluid Dynamics Conference, 1999, p. 3259.
- [12] A. Harten, B. Engquist, S. Osher, S. R. Chakravarthy, Uniformly high order accurate essentially non-oscillatory schemes, III, *Journal of computational physics* 71 (2) (1987) 231–303.
- [13] R. Abgrall, On essentially non-oscillatory schemes on unstructured meshes: analysis and implementation, *Journal of Computational Physics* 114 (1) (1994) 45–58.
- [14] C. F. Ollivier-Gooch, Quasi-ENO schemes for unstructured meshes based on unlimited data-dependent least-squares reconstruction, *Journal of Computational Physics* 133 (1) (1997) 6–17.
- [15] X.-D. Liu, S. Osher, T. Chan, Weighted essentially non-oscillatory schemes, *Journal of computational physics* 115 (1) (1994) 200–212.
- [16] C. Hu, C.-W. Shu, Weighted essentially non-oscillatory schemes on triangular meshes, *Journal of Computational Physics* 150 (1) (1999) 97–127.
- [17] O. Friedrich, Weighted essentially non-oscillatory schemes for the interpolation of mean values on unstructured grids, *Journal of computational physics* 144 (1) (1998) 194–212.
- [18] W. H. Reed, T. Hill, Triangular mesh methods for the neutron transport equation, Los Alamos Report LA-UR-73-479.
- [19] B. Cockburn, C.-W. Shu, The local discontinuous Galerkin method for time-dependent convection-diffusion systems, *SIAM Journal on Numerical Analysis* 35 (6) (1998) 2440–2463.
- [20] J. Peraire, P.-O. Persson, The compact discontinuous Galerkin (CDG) method for elliptic problems, *SIAM Journal on Scientific Computing* 30 (4) (2008) 1806–1824.
- [21] B. Cockburn, J. Gopalakrishnan, R. Lazarov, Unified hybridization of discontinuous Galerkin, mixed, and continuous Galerkin methods for second order elliptic problems, *SIAM Journal on Numerical Analysis* 47 (2) (2009) 1319–1365.

- [22] H. T. Huynh, A flux reconstruction approach to high-order schemes including discontinuous Galerkin methods, AIAA paper 4079 (2007) 2007.
- [23] H. T. Huynh, A reconstruction approach to high-order schemes including discontinuous Galerkin for diffusion, AIAA paper 403 (2009) 2009.
- [24] D. A. Kopriva, J. H. Kolas, A conservative staggered-grid Chebyshev multidomain method for compressible flows, *Journal of computational physics* 125 (1) (1996) 244–261.
- [25] H. T. Huynh, Flux reconstruction / correction procedure via reconstruction, a unified approach to high-order accuracy, 37th Advanced VKI CFD Lecture Series: Recent Developments in Higher Order Methods and Industrial Application in Aeronautics.
- [26] Y. Allaneau, A. Jameson, Connections between the filtered discontinuous Galerkin method and the flux reconstruction approach to high order discretizations, *Computer Methods in Applied Mechanics and Engineering* 200 (49) (2011) 3628–3636.
- [27] D. De Grazia, G. Mengaldo, D. Moxey, P. Vincent, S. Sherwin, Connections between the discontinuous Galerkin method and high-order flux reconstruction schemes, *International journal for numerical methods in fluids* 75 (12) (2014) 860–877.
- [28] P. Zwanenburg, S. Nadarajah, Equivalence between the energy stable flux reconstruction and filtered discontinuous Galerkin schemes, *Journal of Computational Physics* 306 (2016) 343–369.
- [29] M. R. López, A. Sheshadri, J. R. Bull, T. D. Economou, J. Romero, J. E. Watkins, D. M. Williams, F. Palacios, A. Jameson, D. E. Manosalvas, Verification and validation of HiFiLES: a high-order LES unstructured solver on multi-GPU platforms, in: 32nd AIAA Applied Aerodynamics Conference, 2014, p. 3168.
- [30] P. Vincent, F. D. Witherden, A. M. Farrington, G. Ntemos, B. C. Vermeire, J. S. Park, A. S. Iyer, PyFR: Next-generation high-order computational fluid dynamics on many-core hardware, in: 22nd AIAA Computational Fluid Dynamics Conference, 2015, p. 3050.
- [31] F. D. Witherden, B. C. Vermeire, P. E. Vincent, Heterogeneous computing on mixed unstructured grids with PyFR, *Computers & Fluids* 120 (2015) 173–186.
- [32] H. T. Huynh, High-order methods including discontinuous Galerkin by reconstructions on triangular meshes, AIAA Paper 44 (2011) 2011.
- [33] P. Castonguay, P. E. Vincent, A. Jameson, A new class of high-order energy stable flux reconstruction schemes for triangular elements, *Journal of Scientific Computing* 51 (1) (2011) 224–256.
- [34] P. E. Vincent, P. Castonguay, A. Jameson, A new class of high-order energy stable flux reconstruction schemes, *Journal of Scientific Computing* 47 (1) (2011) 50–72.
- [35] A. Jameson, A proof of the stability of the spectral difference method for all orders of accuracy, *Journal of Scientific Computing* 45 (1-3) (2010) 348–358.
- [36] P. Castonguay, D. M. Williams, P. E. Vincent, M. Lopez, A. Jameson, On the development of a high-order, multi-GPU enabled, compressible viscous flow solver for mixed unstructured grids, AIAA paper 3229 (2011) 2011.
- [37] A. Lani, N. Villedie, K. Bensassi, L. Koloszar, M. Vymazal, S. M. Yalim, M. Panesi, COOLFluid: an open computational platform for multi-physics simulation and research, in: 21st AIAA Computational Fluid Dynamics Conference, 2013, p. 2589.
- [38] A. Lani, T. Quintino, D. Kimpe, H. Deconinck, S. Vandewalle, S. Poedts, Reusable object-oriented solutions for numerical simulation of PDEs in a high performance environment, *Scientific Programming* 14 (2) (2006) 111–139.
- [39] A. Lani, T. Quintino, D. Kimpe, H. Deconinck, S. Vandewalle, S. Poedts, The COOLFluid framework: Design solutions for high-performance object oriented scientific computing software, in: V. S. Sunderan, G. D. van Albada, P. M. A. Sloot, J. J. Dongarra (Eds.), *Computational Science ICCS 2005*, Vol. 1 of LNCS 3514, Emory University, Springer, Atlanta, GA, USA, 2005, pp. 281–286.
- [40] F. Bassi, S. Rebay, A high-order accurate discontinuous finite element method for the numerical solution of the compressible Navier-Stokes equations, *Journal of computational physics* 131 (2) (1997) 267–279.
- [41] P. Castonguay, P. Vincent, A. Jameson, Application of high-order energy stable flux reconstruction schemes to the Euler equations, in: 49th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition, 2011, p. 686.
- [42] P. L. Roe, Approximate riemann solvers, parameter vectors, and difference schemes, *Journal of computational physics* 43 (2) (1981) 357–372.
- [43] A. Sheshadri, A. Jameson, Shock detection and capturing methods for high order discontinuous-Galerkin finite element methods, in: 32nd AIAA Applied Aerodynamics Conference, 2014, p. 2688.
- [44] J. S. Park, T. K. Chang, C. Kim, Higher-order multi-dimensional limiting strategy for correction procedure via reconstruction, in: 52nd Aerospace Sciences Meeting, 2014, p. 0772.
- [45] J. Du, C.-W. Shu, M. Zhang, A simple weighted essentially non-oscillatory limiter for the correction procedure via reconstruction (CPR) framework, *Applied Numerical Mathematics* 95 (2015) 173–198.
- [46] M. Yu, F. X. Giraldo, M. Peng, Z. J. Wang, Localized artificial viscosity stabilization of discontinuous Galerkin methods for nonhydrostatic mesoscale atmospheric modeling, *Monthly Weather Review* 143 (12) (2015) 4823–4845.
- [47] Y. Li, Z. J. Wang, A convergent and accuracy preserving limiter for the FR/CPR method, in: 55th AIAA Aerospace Sciences Meeting, 2017, p. 0756.
- [48] P.-O. Persson, J. Peraire, Sub-cell shock capturing for discontinuous Galerkin methods, in: 44th AIAA Aerospace Sciences Meeting and Exhibit, 2006, p. 112.
- [49] J. S. Park, M. Yu, C. Kim, Z. Wang, Comparative study of shock-capturing methods for high-order CPR: MLP and artificial viscosity, *Proceedings of the 8th International Conference on CFD, ICCFD8-2014-0067*.
- [50] X. Zhang, On positivity-preserving high order discontinuous Galerkin schemes for compressible Navier-Stokes equations, *Journal of Computational Physics* 328 (2017) 301–343.
- [51] D. Knight, O. Chazot, J. Austin, M. A. Badr, G. Candler, B. Celik, D. de Rosa, R. Donelli, J. Komives, A. Lani, et al., Assessment of predictive capabilities for aerodynamic heating in hypersonic flow, *Progress in Aerospace Sciences* 90 (2017) 39–53.
- [52] M. Panesi, A. Lani, T. Magin, J. Molnar, O. Chazot, H. Deconinck, Numerical investigation of the non equilibrium shock-layer around the EXPERT vehicle, in: 18th AIAA Computational Fluid Dynamics Conference, 2007, p. 4317.
- [53] G. Degrez, A. Lani, M. Panesi, O. Chazot, H. Deconinck, Modelling of high-enthalpy, high-mach number flows, *J. Phys. D: App. Phys* 41.
- [54] D. Knight, J. Longo, D. Drikakis, D. Gaitonde, A. L. et al., Assessment of CFD capability for prediction of hypersonic shock interactions, *Prog. Aerosp. Sci.* 48-49 (2012) 8–26.
- [55] A. Munafo, A. Lani, A. Bultel, M. Panesi, Modeling of non-equilibrium phenomena in expanding flows by means of a collisional-radiative

- model, *Physics of Plasma* 20 (7).
- [56] M. Panesi, A. Lani, Collisional radiative coarse-grain model for ionization in air, *Physics of Fluids* 25 (2013) 057101.
 - [57] A. Lani, P. Duarte Santos, A. Sanna, An efficient Monte Carlo method for radiation transport in aerothermodynamic simulations, in: 44th AIAA Thermophysics Conference, 2013, p. 2893.
 - [58] A. Lani, An object oriented and high performance platform for aerothermodynamics simulation, Ph.D. thesis, Université Libre de Bruxelles (2008).
 - [59] T. Quintino, A component environment for high-performance scientific computing: design and implementation, Ph.D. thesis, Ph. D. Thesis, KU Leuven, von Karman Institute for Fluid Dynamics, Leuven (2008).
 - [60] B. Van Leer, J. L. Thomas, P. L. Roe, R. W. Newsome, A comparison of numerical flux formulas for the Euler and Navier-Stokes equations, AIAA Paper.
 - [61] V. V. Rusanov, Calculation of interaction of non-steady shock waves with obstacles, *J. Comput. Math. Phys. USSR* 1 (1961) 267–279.
 - [62] M.-S. Liou, A sequel to AUSM: AUSM+, *Journal of computational Physics* 129 (2) (1996) 364–382.
 - [63] M.-S. Liou, A sequel to AUSM, part II: AUSM+-up for all speeds, *Journal of Computational Physics* 214 (1) (2006) 137–170.
 - [64] A. Lani, T. Quintino, Design techniques for high performance multi-physics simulations, VKI Lecture Series on High Performance Computing of Industrial Flows.
 - [65] D. A. Kopriva, Metric identities and the discontinuous spectral element method on curvilinear meshes, *Journal of Scientific Computing* 26 (3) (2006) 301.
 - [66] Y. Saad, M. H. Schultz, Gmres: A generalized minimal residual algorithm for solving nonsymmetric linear systems, *SIAM Journal on scientific and statistical computing* 7 (3) (1986) 856–869.
 - [67] Z. J. Wang, K. Fidkowski, R. Abgrall, F. Bassi, D. Caraeni, A. Cary, H. Deconinck, R. Hartmann, K. Hillewaert, H. T. Huynh, et al., High-order CFD methods: current status and perspective, *International Journal for Numerical Methods in Fluids* 72 (8) (2013) 811–845.
 - [68] L. Krivodonova, M. Berger, High-order accurate implementation of solid wall boundary conditions in curved geometries, *Journal of computational physics* 211 (2) (2006) 492–512.
 - [69] F. Bassi, S. Rebay, High-order accurate discontinuous finite element solution of the 2D Euler equations, *Journal of computational physics* 138 (2) (1997) 251–285.
 - [70] G. H. Golub, J. H. Welsch, Calculation of Gauss quadrature rules, *Mathematics of computation* 23 (106) (1969) 221–230.
 - [71] X. Luo, M. S. Shephard, J.-F. Remacle, The influence of geometric approximation on the accuracy of high order methods, Rensselaer SCOREC report 1.
 - [72] K. Van den Abele, Development of high-order accurate schemes for unstructured grids, Ph.D. thesis, Vrije Universiteit Brussel (2009).
 - [73] D. M. Williams, P. Castonguay, P. E. Vincent, A. Jameson, An extension of energy stable flux reconstruction to unsteady, non-linear, viscous problems on mixed grids, in: 20th AIAA Computational Fluid Dynamics Conference, 2011, pp. 27–30.
 - [74] A. Jameson, P. E. Vincent, P. Castonguay, On the non-linear stability of flux reconstruction schemes, *Journal of Scientific Computing* 50 (2) (2012) 434–445.
 - [75] C. E. Tansley, D. P. Marshall, Flow past a cylinder on a β plane, with application to Gulf Stream separation and the Antarctic Circumpolar Current, *Journal of Physical Oceanography* 31 (11) (2001) 3274–3283.
 - [76] A. Roshko, On the development of turbulent wakes from vortex streets, National Advisory Committee for Aeronautics.
 - [77] M. S. Bloor, The transition to turbulence in the wake of a circular cylinder, *Journal of Fluid Mechanics* 19 (02) (1964) 290–304.
 - [78] C. Williamson, The existence of two stages in the transition to three-dimensionality of a cylinder wake, *The Physics of fluids* 31 (11) (1988) 3165–3168.
 - [79] P. Gnoffo, J. White, Computational aerothermodynamic simulation issues on unstructured grids, in: 37th AIAA Thermophysics Conference, 2004, p. 2371.
 - [80] G. Candler, I. Nompelis, M.-C. Druguet, M. Holden, T. Wadhams, I. Boyd, W.-L. Wang, CFD validation for hypersonic flight-hypersonic double-cone flow simulations, in: 40th AIAA Aerospace Sciences Meeting & Exhibit, 2002, p. 581.
 - [81] A. Lani, Development of an object oriented framework for PDE solvers on unstructured grids, Travail réalisé pour obtenir le DEA en sciences appliquées, Université Libre de Bruxelles, Faculté des Sciences Appliquées (Sep 2003).
 - [82] S. Balay, S. Abhyankar, M. F. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, A. Dener, V. Eijkhout, W. D. Gropp, D. Kaushik, M. G. Knepley, D. A. May, L. C. McInnes, R. T. Mills, T. Munson, K. Rupp, P. Sanan, B. F. Smith, S. Zampini, H. Zhang, H. Zhang, PETSc Web page, <http://www.mcs.anl.gov/petsc> (2019).
URL <http://www.mcs.anl.gov/petsc>