Eleventh International Conference on Computational Fluid Dynamics (ICCFD11), Maui, HI, USA, July 11-15, 2022

ICCFD11-2022-0405

Assessment of a high-order implicit residual smoothing time scheme for multiblock curvilinear meshes

A. Bienner*, X. Gloerfelt* and P. Cinnella**
Corresponding author: xavier.gloerfelt@ensam.eu

* Arts et Métiers, DynFluid Laboratory, 151 bd. de l'Hôpital, 75013 Paris, France.

** Sorbonne University, Jean Le Rond D'Alembert Institute, 4 place Jussieu, 75005 Paris, France.

Abstract: In direct and large eddy simulations, very small space steps are used close to the solid walls in order to resolve the boundary-layer structures. Due to the restrictive CFL stability criteria of explicit time-stepping schemes, the maximum allowable time step is also very small, leading to high computational costs, notably for converging flow statistics. The use of an implicit integration scheme may overcome this limitation at the price of an increased computional cost per step. Furthermore, the most commonly used fully implicit schemes induce higher errors due to the necessary approximations and bad dispersion and dissipation properties. As a compromise, a fourth-order implicit residual smoothing scheme (IRS4), succesfully validated for a finite volume solver in [1, 2], has been implemented in a multiblock high-order finite-difference solver. For moderate CFL numbers, a similar accuracy as the explicit method is obtained with substantial savings in terms of computational time.

Keywords: High-order Numerical Algorithms, Implicit Time Advancement, Residual Smoothing.

1 Introduction

High-order numerical methods are being implemented for direct numerical simulations (DNS) and large-eddy simulations (LES) of turbulent flows and aeroacoustics. In order to preserve the numerical accuracy and due to time scales comparable to the spatial scales, the Courant-Friedrichs-Lewy (CFL) number is close to unity and generally explicit time integration schemes, such as Runge–Kutta or Adams-Bashforth schemes, are used. The limit is then dictated by the physics rather than the numerical stability. However, in the presence of solid boundaries, thin turbulent boundary layers develop on the walls which require the use of very small mesh sizes. In such a situation, stability constraints of explicit schemes impose time steps much smaller than the characteristic time of viscous-dominated turbulent structures near the wall, whereas the physics would instead correspond to CFL of the order of 10. In the incompressible regime, if the stability limit is dictated by a viscous criterion, an implicit scheme can remove the limit at the price of the solution of a linear systems. The cost being moderate the algorithm remains efficient. For compressible Navier-Stokes equations, due to the hyperbolicity, implicit method necessitates the solution of nonlinear systems. If a sufficiently accurate time integration is chosen, such as implicit Runge–Kutta schemes or linear multistep methods (Adams-Moulton or Backward Differentiation Formula schemes), the overcost may lose the benefits of increasing the time step beyond the CFL limit.

A way to relax the severe stability limitations of explicit schemes without going into the trouble of fully implicit ones, is the use of Implicit Residual Smoothing (IRS), a technique originally introduced to accelerate the convergence to steady solutions. In fact, as shown in [1, 2], this technique can be extended to high order and can simulate unsteady flow at CFL between 5 and 10 accurately with a moderate overcost. IRS may be one order of magnitude cheaper than a fully implicit method where the nonlinear system is solved to strict tolerance (even if absolute timing is always a matter of controverse). In the context of compressible flow simulations, IRS was introduced at a time where computer power just allowed the computation of steady

solutions. A relatively accurate (second-order) implicit phase was constructed by Lerat and coworkers [3, 4, 5, 6] by simplifying a Lax-Wendroff-type implicit phase. Based on second-order differencing, it acts as a Laplacian smoother to accelerate convergence. It has great similarity with the implicit phase introduced later by Jameson and coworkers [7, 8] to smooth/filter the solution. IRS was used in several computational fluid dynamics (CFD) codes to speed-up convergence of steady Euler and Navier-Stokes calculations in conjunction with multigrid algorithm [9, 10, 11, 12, 13]. A fourth-order accurate version was introduced in [1] and applied to the LES and DNS of selected geometrically simple flow configurations. A finite-volume multi-domain formulation was discussed in [2] and shown to be efficient for LES in turbomachinery in conjunction with a third-order accurate spatial scheme. The fourth-order IRS relies on a bi-Laplacian filter to damp high-frequency modes of the residual, which leads to the solution of pentadiagonal systems for each space direction and Runge–Kutta stage. Thanks to the efficient inversion of scalar pentadiagonal matrices, the extra computational cost associated with the implicit operator was shown to remain much lower than standard implicit schemes at least for the considered configurations. Nevertheless, its multi-domain implementation requires simplifications of the implicit matrix that may hinder robustness. Furthermore, the numerical dissipation properties of the implicit scheme are also modified and may lead to reduced numerical stability in conjunction with higher-order schemes, such as the ninth order scheme considered in this study.

The objective of the present work is manyfold: i) first, the IRS is extended to a curvilinear finite-difference formulation, more suitable than the finite-volume formulation to achieve high-order accuracy in space; ii) a detailed study of the dissipation is carried out, and a filtering strategy is introduced to ensure the correct damping of small scales at high CFL numbers; iii) finally, the numerical properties of boundary and interface treatment, which is critical for DNS and LES on massively parallel computers, is presented.

The paper is organized as follows: Section 1 presents the governing equations and the time and space discretization schemes used in this study. Section 3 gives the principles and properties of IRS implicit phase, notably in the Fourier space and explains its implementation in a finite-difference code with curvilinear transformation. The problem of boundary and interface treatments are detailed in Section 4. Section 5 give some numerical applications showing the efficiency and the interest of the current time implicitation strategy.

2 Baseline numerical methods

2.1 Governing equations

The governing equations are the compressible Navier-Stokes equations written for a curvilinear domain by using a coordinate transform. The physical space (x,y) is mapped into a Cartesian regular computational space (ξ,η), and the third direction, which corresponds in the following to the spanwise direction z, is left unchanged. By denoting (u,v,w) the velocity components, ρ the density, p the pressure and E the total specific energy, the set of equations for the unknown vector $\mathbf{U} = (\rho, \rho u, \rho v, \rho w, \rho E)^T$ is given by:

$$\frac{\partial \mathbf{U}}{\partial t} + \frac{1}{J}\left(\frac{\partial \mathbf{F}_c}{\partial \xi} + \frac{\partial \mathbf{G}_c}{\partial \eta}\right) + \frac{\partial \mathbf{H}}{\partial z} = 0 \tag{1}$$

with $J = x_\xi y_\eta - y_\xi x_\eta$ the Jacobian of the coordinate transformation. The fluxes $\mathbf{F}_c$, $\mathbf{G}_c$, $\mathbf{H}$ are the sum of the inviscid (superscript e) and visco-thermal fluxes (superscript v) given by

$$\begin{aligned}
\mathbf{F}_c^e &= \begin{pmatrix} \rho J\Theta_\xi \\ \rho u J\Theta_\xi + p y_\eta \\ \rho v J\Theta_\xi - p x_\eta \\ \rho w J\Theta_\xi \\ (\rho E + p) J\Theta_\xi \end{pmatrix}, \quad \mathbf{F}_c^v = \begin{pmatrix} 0 \\ \tau_{xx} y_\eta - \tau_{xy} x_\eta \\ \tau_{xy} y_\eta - \tau_{yy} x_\eta \\ \tau_{xz} y_\eta - \tau_{yz} x_\eta \\ (\mathbf{u}\cdot\boldsymbol{\tau}_x) y_\eta - (\mathbf{u}\cdot\boldsymbol{\tau}_y) x_\eta - (q_x y_\eta - q_y x_\eta) \end{pmatrix}, \\
\mathbf{G}_c^e &= \begin{pmatrix} \rho J\Theta_\eta \\ \rho u J\Theta_\eta - p y_\xi \\ \rho v J\Theta_\eta + p x_\xi \\ \rho w J\Theta_\eta \\ (\rho E + p) J\Theta_\eta \end{pmatrix}, \quad \mathbf{G}_c^v = \begin{pmatrix} 0 \\ \tau_{xy} x_\xi - \tau_{xx} y_\xi \\ \tau_{yy} x_\xi - \tau_{xy} y_\xi \\ \tau_{yz} x_\xi - \tau_{xz} y_\xi \\ (\mathbf{u}\cdot\boldsymbol{\tau}_y) x_\xi - (\mathbf{u}\cdot\boldsymbol{\tau}_x) y_\xi - (q_y x_\xi - q_x y_\xi) \end{pmatrix},
\end{aligned} \tag{2}$$

$$\mathbf{H}^e = \begin{pmatrix} \rho w \\ \rho uw \\ \rho vw \\ \rho w^2 + p \\ (\rho E + p)w \end{pmatrix}, \quad \mathbf{H}^v = \begin{pmatrix} 0 \\ \tau_{xz} \\ \tau_{yz} \\ \tau_{zz} \\ \mathbf{u} \cdot \boldsymbol{\tau}_z - q_z \end{pmatrix}$$

with $\mathbf{u} = (u, v, w)^T$, $\mathbf{q} = (q_x, q_y, q_z)^T$ the heat flux, $\boldsymbol{\tau} = (\boldsymbol{\tau}_x, \boldsymbol{\tau}_y, \boldsymbol{\tau}_z)^T$ the viscous stress tensor, and we have noted $\boldsymbol{\tau}_x = (\tau_{xx}, \tau_{xy}, \tau_{xz})^T$, $\boldsymbol{\tau}_y = (\tau_{xy}, \tau_{yy}, \tau_{yz})^T$ and $\boldsymbol{\tau}_z = (\tau_{xz}, \tau_{yz}, \tau_{zz})^T$. The contravariant velocities are defined as:

$$\Theta_\xi = \mathbf{u} \cdot \nabla\xi = \frac{1}{J}(uy_\eta - vx_\eta) \quad \text{and} \quad \Theta_\eta = \mathbf{u} \cdot \nabla\eta = \frac{1}{J}(vx_\xi - uy_\xi) \tag{3}$$

The specific total energy is $E = p/[(\gamma - 1)\rho] + (u^2 + v^2 + w^2)/2$ for an ideal gas satisfying $p = \rho r T$, where T the temperature, r the gas constant and γ the ratio of specific heats. The tensor $\boldsymbol{\tau}$ follows the Newtonian fluid constitutive relation. The dynamic viscosity μ is approximated with Sutherland's law and the heat flux components are modelled with Fourier's law,.

2.2 Base space and time discretizations

Central differencing is used to approximate the flux derivative. A spatial derivative of a quantity u on a uniform mesh of size Δx by means of a tenth-order accurate central scheme using an eleven-point stencil:

$$\left.\frac{\partial u}{\partial x}\right|_j = \frac{1}{\Delta x} \sum_{l=-5}^{5} a_l \, u\left((j + l)\Delta x\right) \tag{4}$$

The baseline time-stepping procedure is the explicit low-storage Runge–Kutta (RK) of Jameson *et al.* [14], which may be written as:

$$\begin{cases} \mathbf{U}^{(0)} = \mathbf{U}^n \\ \Delta\mathbf{U}^{(k)} = -\alpha_k \Delta t \mathcal{R}(\mathbf{U}^{(k-1)}), \quad k = 1, ..., s \\ \mathbf{U}^{n+1} = \mathbf{U}^{(s)} \end{cases} \tag{5}$$

where $\mathbf{U}^n$ is the numerical solution at time $n\Delta t$, $\Delta\mathbf{U}^{(k)} = \mathbf{U}^{(k)} - \mathbf{U}^{(0)}$ is the solution increment at the k-th RK stage, s=4 is the number of stages and α_k are the scheme coefficients (α_1=1/4, α_2=1/3, α_3=1/2, α_4=1). Its formal accuracy is four for linear equations and reduces down to second order for nonlinear problems.

All the numerical schemes are implemented within the MUSICAA code (MUltiblock Solver In Computational Aerodynamics and Aeroacoustics), developed at DynFluid laboratory, which is a high-fidelity finite-difference (FD) solver written in Fortran and parallelized with MPI library. Various spatial and temporal schemes are available but we restrict the attention for the current developments to the tenth-order FD and four-stage RK scheme. In the following applications, the stencil size is reduced down to second order near physical boundaries and we will use fourth-order central differencing for viscous fluxes (see § 4 for details about boundary conditions).

2.3 Numerical dissipation

Central differencing methods require the addition of some form of artificial dissipation. Even starting with a smooth solution, nonlinearity can excite odd-even decoupling and dissipation or filtering can be used to control grid-to-grid oscillations. The latter are preferentially generated near the approximate physical boundary conditions and can arise due to the successive application of the first-order derivative operator in computing the derivatives of viscous fluxes. The problem is enhanced in the presence of sharp flow discontinuities, such as shocks, and in the context of large-eddy simulations. The unresolved modes introduce an aliasing error and have to be dissipated. The particular nature, amount and form of the dissipation has a strong impact on the solution accuracy [15, 16]. In Ref.[1], a scalar artificial dissipation (AD) was used (we omit here the shock-capturing term for simplicity). The approximation of flux derivatives is supplemented by a ninth-order dissipation term:

$$\mathcal{D}_j^{AD} = \frac{1}{1260\Delta x} \delta\left(\lambda^e \delta^9 U_j\right) \tag{6}$$

where δ is the classical difference operator over one cell:

$$\delta(\bullet)_{j+\frac{1}{2}} = (\bullet)_{j+1} - (\bullet)_j \tag{7}$$

and λ^e is the spectral radius of the inviscid flux Jacobian. The coefficient 1/1260 is used by analogy with a ninth-order MUSCL scheme [1].

In previous studies [15, 16], we also used selective filtering (SF). A filtered quantity u_j^* is computed by using a tenth-order centered filter built on an eleven-point stencil [17, 18] :

$$u_j^* = u_j - \chi \mathcal{D}_j^{SF} \quad \text{with} \quad \mathcal{D}_j^{SF} = \sum_{l=-5}^{5} d_l \, u\left((j+l)\Delta x\right) , \tag{8}$$

with the coefficient χ ranging between 0 and 1. In most applications, a value $\chi \in [0.1\,, 0.2]$ is used, so that a minimal amount of dissipation is introduced. The filter has symmetric coefficients d_l, so that it is nondispersive. The connection with artificial dissipation has been clearly highlighted by Edoh *et al.* [19] (see also [20, 21]), who showed that the performance of explicit AD and SF are similar. Both methods adjust the amount of dissipation to the mesh size Δx. One pitfall of the filtering approach is that it does not take into account the time step Δt, and for very small time steps (or CFL numbers), the repeated application of the filter can lead to an overdissipation. A simple trick proposed in [19] to overcome this artefact is to replace χ by $\min\{\chi, \mathrm{CFL}\}$, which is done in the present implementation. We will see that selective filtering independent on Δt has also some advantages in the context of IRS time advancement.

3 Time integration using high-order Implicit Residual Smoothing

3.1 Principle of Implicit Residual Smoothing

In the original developments by Lerat and coworkers [3, 4, 5, 6], an implicit phase was proposed for a general class of three-point schemes using an approximate Lax-Wendroff-like implicit operator. Starting from the 2D Euler equations, written as:

$$W_t + f(W)_x + g(W)_y = 0$$

and denoting $A = \partial f/\partial w$ and $B = \partial g/\partial w$ the Jacobian matrices, the complete implicit phase is:

$$\Delta W + \theta \frac{\Delta t^2}{2} \left\{ \left[A^2(\Delta W)_x + AB(\Delta W)_y \right]_x + \left[BA(\Delta W)_x + B^2(\Delta W)_y \right]_y \right\} = \Delta W^{expl}$$

where ΔW is the solution increment and ΔW^{expl} represents the explicit phase. It is unconditionally stable for $\theta \leq -1/2$. A first simplification is the suppression of cross-derivative terms:

$$\Delta W + \theta \frac{\Delta t^2}{2} \left\{ \left[A^2(\Delta W)_x \right]_x + \left[B^2(\Delta W)_y \right]_y \right\} = \Delta W^{expl}$$

An approximate factorization per spatial directions is then used, yielding:

$$\Delta W^* + \theta \frac{\Delta t^2}{2} \left[A^2(\Delta W^*)_x \right]_x = \Delta W^{expl} \quad \text{then} \quad \Delta W + \theta \frac{\Delta t^2}{2} \left[B^2(\Delta W)_y \right]_y = \Delta W^*$$

It is further simplified by replacing the Jacobian matrices by their spectral radius:

$$\Delta W^* + \theta \frac{\Delta t^2}{2} \left[\rho(A)^2(\Delta W^*)_x \right]_x = \Delta W^{expl} \quad \text{then} \quad \Delta W + \theta \frac{\Delta t^2}{2} \left[\rho(B)^2(\Delta W)_y \right]_y = \Delta W^*$$

After discretization by second-order central differences, the implicit operator $\mathcal{J}$ for each direction reads:

$$\mathcal{J} = 1 - \theta_2 \left(\frac{\Delta t}{\Delta x} \right)^2 \delta(\lambda^{e\,2} \delta) \tag{9}$$

where δ is the differencing operator (7), λ^e denotes the spectral radius of the flux Jacobians in each direction and θ_2 is the value of the coefficient for the second-order scheme. A tridiagonal system has to be solved for each space dimension.

A similar implicit operator was used by Jameson & Baker [7] to increase the CFL number of an explicit Runge–Kutta time-stepping scheme for steady problems. The so-called "implicit residual smoothing" is applied at each RK stages as:

$$\left\{ \begin{array}{l} W^{(0)} = W^n \\ \mathcal{J}\Delta W^{(k)} = -\alpha_k \Delta t \mathcal{R}(W^{(k-1)}), \quad k = 1, ..., s \\ W^{n+1} = W^{(s)} \end{array} \right. \tag{10}$$

where the implicit operator $\mathcal{J}$ of Jameson & Baker (JB) reads:

$$\mathcal{J}_{JB} = 1 - \varepsilon \left(\frac{\Delta t}{\Delta x} \right)^2 \delta^2 \quad \text{with} \quad \varepsilon \geq \frac{1}{4} \left[\left(\frac{\text{CFL}}{\text{CFL}^{expl}} \right)^2 - 1 \right]$$

which has only first-order accuracy [22], whereas the Lax-Wendroff IRS operator (9) preserves second-order accuracy. The Lax-Wendroff IRS2 was used in combination with a RK time stepping by Cinnella & Lerat [23]. An extension to fourth-order (IRS4) was first proposed by Cinnella and Content [1], by keeping the same form of the operator (9) but replacing the Laplacian smoothing by a bi-Laplacian smoothing:

$$\mathcal{J} = 1 + \theta_4 \left(\frac{\Delta t}{\Delta x} \right)^4 \delta(\lambda^{e4} \delta^3) \,, \tag{11}$$

which implies the resolution of a pentadiagonal system for each space directions. IRS4 is unconditionally stable for $\theta_4 \gtrapprox 0.0023$ (see § 4.3). We use the value $\theta_4 = 0.0025$ in the following. The additional error introduced by the IRS4 operator with respect to the explicit scheme takes the form [1]:

$$-\theta_4 \Delta t^4 \lambda^{e4} \frac{\partial^5 \mathbf{F}^e}{\partial x^5} + \mathcal{O}(\Delta t^4) \tag{12}$$

with $\mathbf{F}^e$ the inviscid flux in the considered direction. Being proportional to an odd derivative, this error has a dispersive nature. As the RK scheme is second-order in general, the additional error coming from the IRS4 is expected to have a negligible impact on the accuracy of the baseline scheme.

3.2 IRS operator for curvilinear coordinates

Using 2D coordinate transform $(\xi, \eta, z) \rightarrow (x, y, z)$ in § 2.1, the multi-dimensional implicit operator is written as:

$$\mathcal{J} = \left(1 + \theta_4 \Delta t^4 \delta_\xi ({\lambda^e_{\xi\eta}}^4 \delta^3_\xi) \right) \left(1 + \theta_4 \Delta t^4 \delta_\eta ({\lambda^e_{\xi\eta}}^4 \delta^3_\eta) \right) \left(1 + \theta_4 \left(\frac{\Delta t}{\Delta z} \right)^4 \delta_z (\lambda^{e4}_z \delta^3_z) \right) \tag{13}$$

noting that the transformed coordinates (ξ, η) correspond to a unitary Cartesian grid ($\Delta\xi = \Delta\eta = 1$). The subscript of the differencing operator δ denotes the grid direction in which it is applied. In the z-direction, the directional spectral radius of the inviscid Jacobian matrix is $\lambda^e_z = |w| + c$, c being the sound speed. Since the transformed directions are coupled, a single spectral radius is used for (ξ, η)-directions (in the same spirit as Pulliam [24] for artificial dissipation):

$$\lambda^e_{\xi\eta} = \sqrt{\Theta^2_\xi + \Theta^2_\eta} + c\sqrt{\nabla\xi^2 + \nabla\eta^2} \,, \tag{14}$$

the contravariant velocities Θ_ξ and Θ_η being defined by (3).

3.3 Interpretation in Fourier space

The main idea of IRS is to stabilize the scheme by smoothing the residuals by means of a Laplacian filter (IRS2) or a bi-Laplacian filter (IRS4). The IRS operator acts on the difference scheme by contracting its support and thus relaxes the restriction on the time step imposed by the CFL condition, which can be easily interpreted using Fourier symbols of the operators.

The amplification factor of the RK4 scheme is found to be:

$$\widehat{g}(\zeta) = 1 + \zeta + \zeta^2/2 + \zeta^3/6 + \zeta^4/24\,, \quad \zeta \in \mathbb{C} \tag{15}$$

The stability region can be represented in the complex plane. The black line in Fig.1 marks the stability boundary ($|\widehat{g}(\zeta)| = 1$). The gray color area denotes the unstable region ($|\widehat{g}(\zeta)| > 1$). Consider the linear scalar advection problem

$$\frac{\partial W}{\partial t} = -a\frac{\partial W}{\partial x} = \mathcal{R}(W) \quad (a > 0) \tag{16}$$

The locus of the FD operator (4) is given by :

$$\begin{aligned} a\Delta x\widehat{\mathcal{R}}(k\Delta x) &= 2i\sum_{l=1}^{5} a_l \sin(lk\Delta x) \\ &= 2i\left(\frac{5}{6}\sin(k\Delta x) - \frac{5}{21}\sin(2k\Delta x) + \frac{5}{84}\sin(3k\Delta x) - \frac{5}{504}\sin(4k\Delta x) + \frac{1}{1260}\sin(5k\Delta x)\right) \\ &= i\sin\left(\frac{k\Delta x}{2}\right)\cos\left(\frac{k\Delta x}{2}\right)\left[1 + \frac{4}{3}\sin^3\left(\frac{k\Delta x}{2}\right) + \frac{16}{15}\sin^5\left(\frac{k\Delta x}{2}\right) + \frac{32}{35}\sin^7\left(\frac{k\Delta x}{2}\right) + \frac{256}{315}\sin^9\left(\frac{k\Delta x}{2}\right)\right] \end{aligned} \tag{17}$$

with $i = \sqrt{-1}$. The space discretization is nondissipative and its locus lies on the imaginary axis.
The Fourier symbol of the artificial dissipation operator (6) is given by:

$$|a|\Delta x\widehat{\mathcal{D}}^{AD}(k\Delta x) = \frac{256}{315}\sin^{10}\left(\frac{k\Delta x}{2}\right) \tag{18}$$

and its locus is on the real axis. In order for the discrete scheme to be stable, the locus of the spatial operator $(\widehat{\mathcal{R}} + \widehat{\mathcal{D}}^{AD})$ must remain within the stability region of the RK integration scheme, as shown in Fig.1a. As the CFL increases, the locus will expand and ultimately exit the stability region, as exemplified in the subsequent subfigures.

The IRS smoothing in Eq.(10) corresponds to a modification of the Fourier symbol of the spatial operator, namely $(\widehat{\mathcal{R}} + \widehat{\mathcal{D}}^{AD})/\widehat{\mathcal{J}}$, where the Fourier symbol of IRS4 reads:

$$\widehat{\mathcal{J}} = 1 + 16\,\theta_4 \mathrm{CFL}^4 \sin^4\left(\frac{k\Delta x}{2}\right) \tag{19}$$

The modified locii for IRS4 are plotted in Fig.1 for CFL=1 to 6. The IRS operator contracts the spatial locus, which is flattened on the imaginary axis due to the dispersive nature of the IRS operator. This also implies that, as the CFL increases, the dissipation error decreases and tends to zero.

3.4 Von Neumann stability

To better understand the role of the numerical dissipation in combination with IRS smoothing, a Von Neumann stability analysis is conducted. In the linear scalar case of a transport equation $W_t + aW_x = 0$ $(a > 0)$, the RK algorithm is developed as :

$$W_j^{n+1} = W_j^n + \sum_{s=1}^{4} \gamma_s \Delta t^s \frac{\partial^s W_j^n}{\partial t^s} + \ldots \quad \text{where} \quad \gamma_s = \prod_{q=4-s+1}^{4} \alpha_q \tag{20}$$

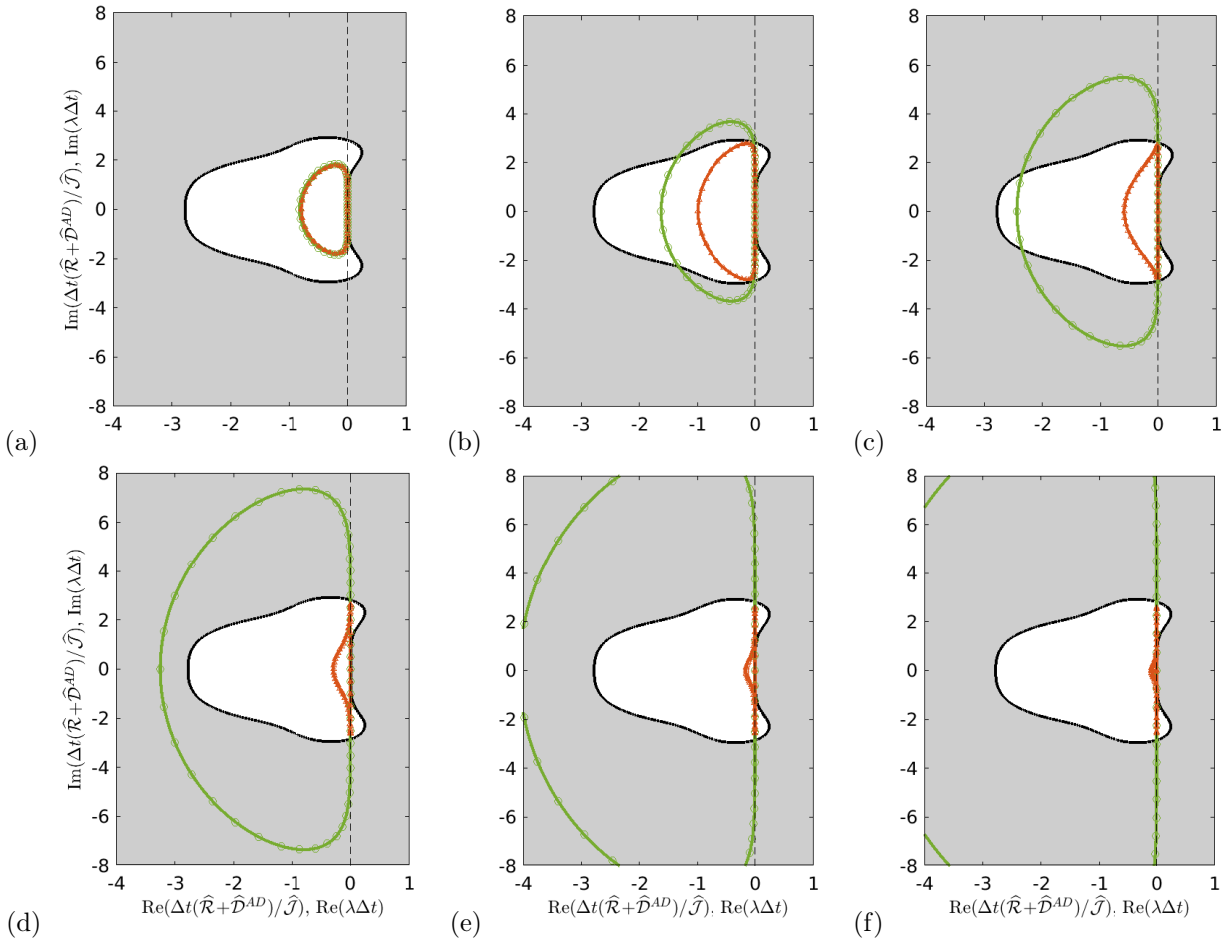

Figure 1: Stability region of the RK4 (white) and locus of the spatial operator with IRS4 smoothing (red) and without (explicit, green), obtained from the Fourier symbols (lines) and the operator's eigenvalues (symbol). From left to right and top to bottom, CFL=1, 2, 3, 4, 5 and 6.

where W_j^n is the solution at grid index j and time iteration n. By applying a Fourier transform to Eq.(20), we write the amplification factor for the discretized equation without numerical dissipation as :

$$\widehat{\mathcal{G}}^0 = \frac{\hat{W}_j^{n+1}}{\hat{W}_j^n} = 1 + \sum_{s=1}^{4} \gamma_s \sigma^s \widehat{\mathcal{R}}^s \tag{21}$$

where $\sigma = a\Delta t/\Delta x$ is the CFL number. Introducing artificial dissipation, the amplification factor becomes:

$$\widehat{\mathcal{G}}^{AD} = 1 + \sum_{s=1}^{4} \gamma_s \sigma^s \left(\widehat{\mathcal{R}} + \widehat{\mathcal{D}}^{AD}\right)^s \tag{22}$$

On the other hand, the filtering operation Eq.(8), applied at the last RK stage, corresponds to a convolution operation in physical space, so that the amplification factor after filtering is multiplied by $1 - \chi\widehat{\mathcal{D}}^{SF}$:

$$\widehat{\mathcal{G}}^{SF}(k\Delta x) = \left(1 + \sum_{s=1}^{4} \gamma_s \sigma^s \widehat{\mathcal{R}}^s\right)\left(1 - \chi\widehat{\mathcal{D}}^{SF}(k\Delta x)\right) \tag{23}$$

where the Fourier symbol of SF is obtained by applying a Fourier transform to Eq.(8) :

$$\begin{aligned}\widehat{\mathcal{D}}^{SF}(k\Delta x) &= d_0 + \sum_{l=1}^{5} 2d_l \cos(lk\Delta x) \\ &= \frac{63}{256} - \frac{105}{256}\cos(k\Delta x) + \frac{15}{64}\cos(2k\Delta x) - \frac{45}{512}\cos(3k\Delta x) + \frac{5}{256}\cos(4k\Delta x) - \frac{1}{512}\cos(5k\Delta x) \\ &= \sin^{10}\left(\frac{k\Delta x}{2}\right) \end{aligned} \tag{24}$$

Note that, for $a = 1$, the Fourier symbol of the tenth-order filter is the same as the one of the ninth-order artificial dissipation when chosing $\chi = 256/315$.

The role of the numerical dissipation is displayed in Fig.2 for CFL numbers between 1 and 15. We first focus on the base scheme without stabilization in Fig.2(a). For the selected value of the IRS coefficient (θ_4=0.0025), it is apparent that the scheme remains always stable. At CFL=1, we observe a damping in the mid-wavenumber range, which depends on the coupled effects of the RK and FD schemes. For higher values of the CFL, the IRS operator acts and a double damping peak is visible, which moves toward low wavenumbers as CFL increases. Damping vanishes at smallest scales ($k\Delta x \approx \pi$), which highlights the need for additional high-wavenumber dissipation. Figure 2(b) shows the amplification factors with artificial dissipation. Since it is added to the spatial operator, its effect is coupled with the temporal scheme, which leads to the non-monotonic behavior for CFL 1 and 2. For high values of CFL, the contractive character of the IRS operator dramatically reduces the dissipation, which could also have been inferred from Fig.1. As a consequence, for CFL $\gtrapprox 10$, almost no dissipation is applied near the grid cut-off ($k\Delta x \approx \pi$), so that grid-to-grid oscillations can remain undamped and pollute the solution or become unstable. In Fig.2(c), the selective filtering is applied with a coefficient $\chi = 256/315$, so that the amount of damping is similar to that of AD without IRS. As expected, the spectral characteristics are a combination of the base scheme and the filter response. In particular, the damping at $k\Delta x \approx \pi$ remains the same independent of the CFL value. The fact that the stabilization is independent of the base scheme gives the greater robustness of RK-IRS implementation, notably for high values of CFL, since the filter will continue to dissipate the modes not resolved by the grid. That is why, in the following, the results are obtained with the base scheme supplemented by the selective filtering.

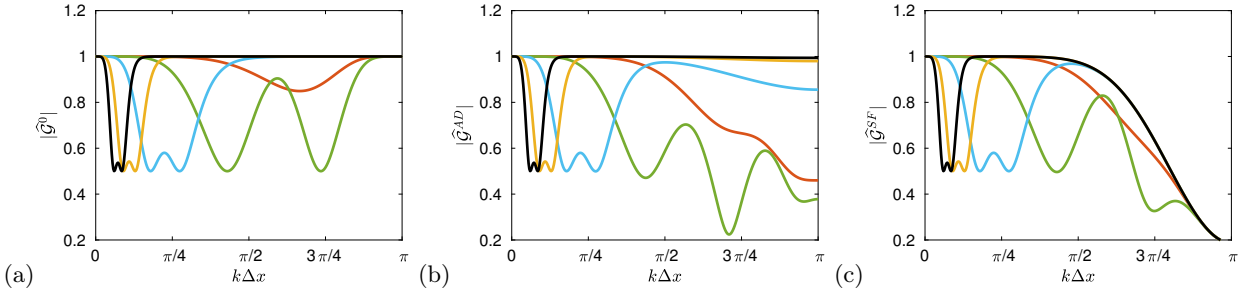

Figure 2: Amplification factors for the centred spatial operator (a), supplemented by artificial dissipation (b) or by selective filtering (c) for increasing CFL numbers, CFL=1 (——), 2 (——), 5 (——), 10 (——), 15 (——).

Since the IRS operator essentially introduces a dispersive error, it is also interesting to look at the relative phase error $\phi + \sigma k\Delta x$ in Fig.3. Since the dissipation term is non dispersive, this quantity is due uniquely to the centred difference operator (4). The striking point in Fig.3(a) is that IRS damages the phase in proportion to the CFL rise. Such a representation would suggest that dispersion errors are unacceptable for CFL greater than 5. However, the analysis would be only valid for a regular grid with constant Δx. In practice, grid points are clustered near wall boundaries, for instance, and high values of the CFL, 5 or greater, are located at these particular points due to the small Δx. In other words, we can also see the spectral properties relative to a fixed time step Δt, which is done in Fig.3(b) and (c) for the amplification factor and the phase respectively by multiplying the abscissae by the CFL number σ. For the advection equation, $\omega = k\,a$, and $\sigma\, k\Delta x = \omega\Delta t$. In this representation, it is clear that the base scheme spectral properties are

preserved in proportion to the CFL number up to $\omega\Delta t = \pi/2$, that is to say if the time signal is discretized by four iterations or more.

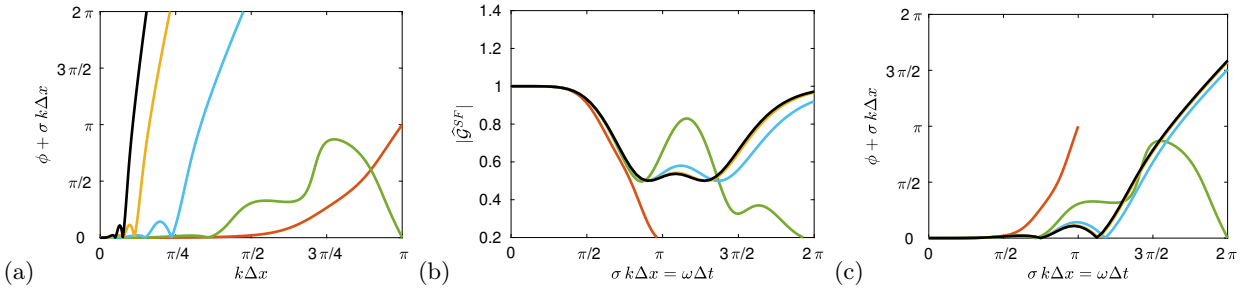

Figure 3: Spectral properties of IRS operator: relative phase error (a), damping factor versus $\omega\Delta t$ (b), and phase error versus $\omega\Delta t$ (c) for increasing CFL numbers, CFL=1 (—), 2 (—), 5 (—), 10 (—), 15 (—).

4 Boundary conditions and multi-domain interface treatment

The preceding analysis is valid for an infinite computational domain or a domain with periodic conditions. The modification of the spatial operator close to boundaries will affect both stability and numerical errors. For this reason, the stability analysis in the following will be carried out with the method of lines [25]. For the linear scalar wave problem (16), spatial discretization on a grid with n points yields

$$\frac{\partial \mathbf{W}_n}{\partial t} = -a\mathbf{R}_{n,n}\mathbf{W}_n$$

where $\mathbf{W}_n$ is the numerical solution vector, of size n and $\mathbf{R}_{n,n}$ is the matrix associated to the spatial operator $\mathcal{R}$. By taking the Fourier transform of this expression, with $\mathbf{W}_n = \hat{\mathbf{W}}_n e^{\Lambda_k t}$, we obtain the following eigenvalue problem :

$$(a\mathbf{R}_{n,n})\hat{\mathbf{W}}_n = -\Lambda_k \hat{\mathbf{W}}_n \tag{25}$$

4.1 Periodic boundary condition

Writing the IRS scheme at the k^{th} RK stage in matrix-vector form

$$\mathbf{J}_{n,n}\Delta\mathbf{W}_n^{(k)} = -\alpha_k \Delta t\, a\mathbf{R}_{n,n}\mathbf{W}_n^{(k-1)}$$

the IRS4 operator leads to the inversion of a pentadiagonal matrix $\mathbf{J}_{n,n}$ per direction at each Runge–Kutta stage. For periodic boundary conditions, the periodic matrix $\mathbf{J}_{n,n}$ reads:

$$\mathbf{J}_{\mathbf{P}n,n} = \begin{pmatrix}
\gamma & \alpha & \beta & 0 & \cdots & \cdots & \cdots & 0 & \beta & \alpha \\
\alpha & \gamma & \alpha & \beta & 0 & \cdots & \cdots & \cdots & 0 & \beta \\
\beta & \alpha & \gamma & \alpha & \beta & 0 & \cdots & \cdots & \vdots & 0 \\
0 & \beta & \alpha & \gamma & \alpha & \beta & 0 & \cdots & \vdots & \vdots \\
\vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\
0 & \cdots & 0 & \beta & \alpha & \gamma & \alpha & \beta & \cdots & \vdots \\
\vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\
0 & \cdots & \cdots & \cdots & 0 & \beta & \alpha & \gamma & \alpha & \beta \\
\beta & 0 & \cdots & \cdots & \cdots & 0 & \beta & \alpha & \gamma & \alpha \\
\alpha & \beta & 0 & \cdots & \cdots & \cdots & 0 & \beta & \alpha & \gamma
\end{pmatrix} \tag{26}$$

with $\beta = \theta_4 \mathrm{CFL}^4$, $\alpha = -4\beta$ and $\gamma = 1 + 6\beta$. The eigenvalue problem (25) then becomes:

$$(a\mathbf{J}_{n,n}^{-1}\mathbf{R}_{n,n})\hat{\mathbf{W}}_n = -\Lambda_k \hat{\mathbf{W}}_n \tag{27}$$

For periodic boundary conditions, as expected, the solutions from the eigenvalue problem (27) with IRS matrix $\mathbf{J}_{\mathbf{P}_{n,n}}$ perfectly match the locus of the Fourier symbol $\widehat{\mathcal{R}}/\widehat{\mathcal{J}}$, as seen in Fig.1 for CFL 1 to 6.

4.2 Physical boundary conditions

Near boundaries of the computational domain it is necessary to modify the IRS matrix. In Ref.[1], the pentadiagonal matrix is simply truncated, which gives matrix $\mathbf{J}_{\mathbf{1}n,n}$ in Eq.(28). To gain some robustness in the FD implementation, we have chosen to reduce IRS order near boundaries. The second to last row of nodes is smoothed with IRS2 operator of Eq.(9) and the one-sided IRS1 operator [12] is used at the boundary points, yielding matrix $\mathbf{J}_{\mathbf{2}n,n}$.

$$\mathbf{J}_{\mathbf{1}n,n} = \begin{pmatrix} \gamma & \alpha & \beta & 0 & \cdots & \cdots & \cdots & 0 \\ \alpha & \gamma & \alpha & \beta & \ddots & \cdots & \cdots & 0 \\ \beta & \alpha & \gamma & \alpha & \beta & \ddots & \cdots & 0 \\ 0 & \beta & \alpha & \gamma & \alpha & \beta & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & 0 \\ 0 & \cdots & \ddots & \beta & \alpha & \gamma & \alpha & \beta \\ 0 & \cdots & \cdots & \ddots & \beta & \alpha & \gamma & \alpha \\ 0 & \cdots & \cdots & \cdots & 0 & \beta & \alpha & \gamma \end{pmatrix}, \quad \mathbf{J}_{\mathbf{2}n,n} = \begin{pmatrix} \gamma_1 & \alpha_1 & 0 & \cdots & \cdots & \cdots & \cdots & 0 \\ \alpha_2 & \gamma_2 & \alpha_2 & 0 & \cdots & \cdots & \cdots & 0 \\ \beta & \alpha & \gamma & \alpha & \beta & 0 & \cdots & 0 \\ 0 & \beta & \alpha & \gamma & \alpha & \beta & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & \beta & \alpha & \gamma & \alpha & \beta \\ 0 & \cdots & \cdots & \cdots & 0 & \alpha_2 & \gamma_2 & \alpha_2 \\ 0 & \cdots & \cdots & \cdots & \cdots & 0 & \alpha_1 & \gamma_1 \end{pmatrix} \tag{28}$$

with coefficients $\alpha_2 = -\theta_2 \mathrm{CFL}^2$ and $\gamma_2 = 1 - 2\alpha_2$ for IRS2, and $\alpha_1 = \pm\theta_1 \mathrm{CFL}$, $\gamma_1 = 1 - \alpha_1$ for IRS1 operator, which is defined as:

$$\begin{cases} \mathcal{J}_{IRS1^+} = 1 - \theta_1 \dfrac{\Delta t}{\Delta x}\delta^+(\lambda^e) & \text{for a left boundary} \\ \\ \mathcal{J}_{IRS1^-} = 1 + \theta_1 \dfrac{\Delta t}{\Delta x}\delta^-(\lambda^e) & \text{for a right boundary} \end{cases} \tag{29}$$

with δ^+, δ^- the upward and backward difference operator respectively. The Fourier symbol of IRS1 and IRS2 are given by:

$$\begin{cases} \widehat{\mathcal{J}}_{IRS1^+} = 1 + 2\,\theta_1 \mathrm{CFL}\left[\sin\left(\dfrac{k\Delta x}{2}\right) - i\cos\left(\dfrac{k\Delta x}{2}\right)\sin\left(\dfrac{k\Delta x}{2}\right)\right] \\ \widehat{\mathcal{J}}_{IRS2} = 1 + 4\,\theta_2 \mathrm{CFL}^2 \sin^2\left(\dfrac{k\Delta x}{2}\right) \end{cases} \tag{30}$$

To determine the coefficients θ ensuring the unconditional stability of the different IRS operators, a numerical search is performed for the 1D scalar problem. The results are plotted in Fig.4. For IRS1, we set the value $\theta_1 = 0.36$. A theoretical value for IRS2 of $1/16$ was obtained analytically in [26, 23] and is used in following, even if the numerical value is always lower than 0.04. For IRS4, the distribution of θ_4 as function of CFL follows the same trend as IRS2 with a peak around CFL=2, where the contractive effect of IRS operator is not yet large. We set the value $\theta_4 = 0.0025$ in the applications.

For instance, for wall boundary conditions, the FD and dissipation stencil is reduced down to second-order for the second to last row. For a no-slip wall, the velocity vector is set to zero ($u = v = w = 0$), a zero normal temperature gradient (adiabatic wall) or a given wall temperature (isothermal wall) are prescribed, and the normal pressure derivative is set to zero. Density is then calculated using the thermal equation of state. For these Dirichlet conditions, the solution increments at the boundary are set equal to zero before

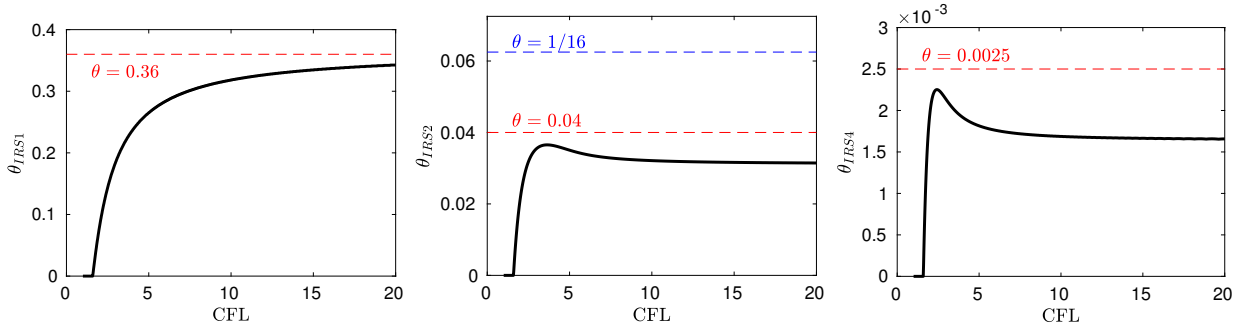

Figure 4: Numerical determination of smoothing parameter θ ensuring unconditional stability for the 1D advection problem: IRS1 (left), IRS2 (center) and IRS4 (right).

applying IRS smoothing.

For a Dirichlet boundary condition, the method of lines is applied with the matrix $\mathbf{J_2}_{n,n}$ and centred schemes with reduced-stencil schemes near boundaries in the spatial operator matrix. The solution increments are imposed at the left boundary so that the corresponding first row and first column are suppressed in the matrices [27]. A simple extrapolation is used at the right boundary. The eigenvalues, plotted in Fig.5 for CFL 2, 5 and 10, show that the locii are significantly modified compared to the old treatment using the truncated matrix $\mathbf{J_1}_{n,n}$. In particular, for the highest CFL value, the use of a reduced stencil IRS operator appears more dissipative, whereas the close-up view in the inset of Fig.5(right) indicates that some eigenvalues can excite weak instabilities with the truncated version. However, using IRS2 and IRS1 coefficients on the borders, all the eigenvalues have a negative real part so that the linear problem has non growing solutions.

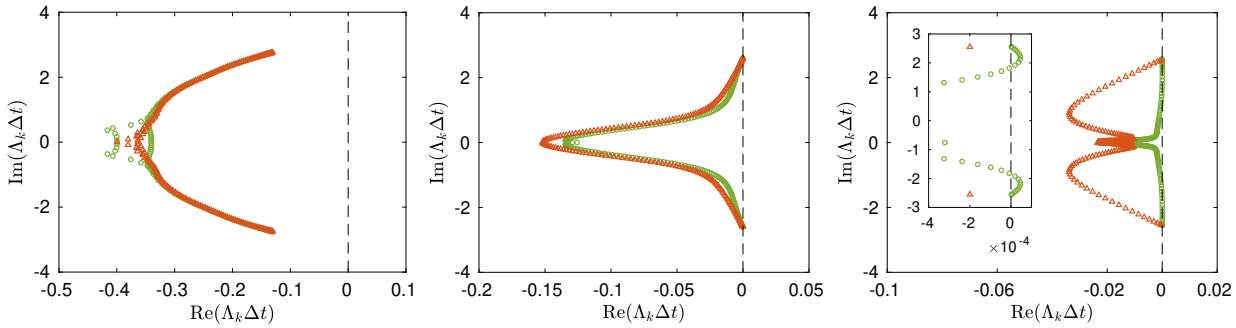

Figure 5: Eigenvalue spectrum with the method of lines for Dirichlet boundary condition using IRS matrix $\mathbf{J_2}_{n,n}$ (△) and $\mathbf{J_1}_{n,n}$ (○): CFL=2 (left), CFL=5 (center) and CFL=10 (right).

4.3 Interface treatment in multi-domain calculations

4.3.1 The different strategies of parallelisation for the linear system resolution

The parallel implementation of IRS is crucial for the efficiency of the method. As we have seen, IRS4 leads to the inversion of a pentadiagonal matrix per mesh direction at each Runge–Kutta step. As the number of linear systems to solve is proportional to the number of implicited directions and to the number of points in that direction, the computational cost can be important. Fortunately, efficient algorithms, such as Thomas' algorithm or cylic reduction, exist to solve small-banded matrices. They are based on a first forward sweep, during which new coefficients are computed, and then a backward substitution to compute the solution. These algorithms are however hard to parallelize efficiently, even if a large literature has tackled this problem, essentially for tridiagonal systems. Strategy based on the redistribution of data between processors [28] or Thomas-pipelined algorithms [29] are better suited for distributed memory architectures.

They lack efficiency due to large communications for the former and idle time for the second. Some methods are based on the divide-and-conquer algorithm, where successive transformations are used to divide the orginal problem into computing cores. The tridiagonal and banded-matrix solvers available in ScaLapack library [30] propose an MPI implementation. We have tested the banded-matrix routines of ScaLapack to solve the pentadiagonal systems but the cost is prohibitive compared to Thomas' algorithm. More complex algorithms, such as Spike, [31] have been proposed to reduce the cost. Even if the scalability has been improved, notably for large band matrices on distributed memory or with GPU processors, the cost remains high for pentadiagonal systems. We have also tested the recent PaScal_tdma library [32], which proposes an optimized parallel cyclic reduction method for tridiagonal matrices. The overcost is more reasonable but the library has to be written for pentadiagonal matrices.

The other possibility to extend IRS to parallel and multiblock configurations is to solve approximate systems on each subdomains by truncating the global system in order to generate local independent linear systems for each processors. Some studies discuss the multi-domain extension for IRS implicitation but are generally restricted to steady problems. Borel and Roux [33] used a Schwarz iterative method for overlapping domain decomposition. They studied the influence of the thickness of the overlap and found an optimal efficiency for an overlapping of L_0=2CFL+1 cells. Lerat and Wu [34] proposed a time-lagging interface condition, where the interface values are fixed at the previous time step. They showed how to obtain a stable and conservative condition. Wu and Zou [35] further analysed this interface treatment and found an optimal overlap L_0=CFL for steady problems. They extended the time-lagging method to unsteady problem by using an overlapping L_0=2CFL+1 points. In Ref.[2], a strategy based on the use of an overlap between blocks was studied by defining halos of ghost points, which is a popular strategy to implement compact FD schemes [36]. It was shown that adding ghost points at the interfaces can limit the error generated by the truncation. This strategy is pursued in the present study.

4.3.2 The ghost-point strategy

Layers of ghost cells are used to make mesh blocks independent and reduce the required number of parallel communications. They are filled using the solution increments computed in the neighboring subdomain at the previous RK stage. In Ref.[1, 2], the IRS operator is simply set equal to the identity for the last two rows, which means that the two outside ghost cells are advanced explicitly in time, while the right-hand sides are communicated from the neighboring block. In the present implementation, the coefficients for the two first/last rows in IRS matrix are those of IRS1 and IRS2 operators. For an interior subdomain, the IRS matrix is thus $J_{2n+ngh,n+ngh}$ as defined in Eq.(28) where n is replaced by $n+ngh$, ngh being the number of ghost cells. Note that the eleven-point stencil base schemes already necessitates five rows of ghost points in each direction, so that we use at least $ngh = 5$. This results in the inversion of a $(n+ngh) \times (n+ngh)$ pentadiagonal matrix per direction on each mesh block, which introduces an overcost. The part due to parallel communications is rather small, and a good parallel scalability is observed for blocks of approximately 50^3 points. The strategy may however suffer from a lack of robustness when using high values of CFL numbers (greater than about 10). To illustrate this point, the scalar advection (16) is solved for a domain $0 < x < 1$ discretized by 600 points. A 1D Gaussian hump, defined as $W(x,0) = \exp(-500(x-0.5)^2)$, is initially located at the middle of the domain and advected periodically to its starting position 10 times at CFL=5. The results are reported in Fig.6 using two solutions: the true periodic solution, *i.e.* IRS matrix is $\mathbf{J}_{\mathbf{p}_{n,n}}$ of Eq.(26) and the approximate ghost-point solution, *i.e.* IRS matrix $\mathbf{J}_{\mathbf{2}_{n+ngh,n+ngh}}$ of Eq.(28). The number of grid points has been chosen so that the periodic hump is correctly advected (some oscillations due to dispersion are visible on its left foot). Then ghost points are used at the left and right boundaries, which can represent the connection with other MPI domains or blocks (note that periodicity is treated as communications by MPI library). Using the nominal number of ghost cells, $ngh = 5$ (first row of Fig.6), a numerical instability is observed after 10 turnovers and the simulation will rapidly blow up. Using an extended number of ghost cells, $ngh = 11$ (bottom row of Fig.6), the solution is in perfect agreement with the periodic one. The corresponding eigenspectrum are plotted on the left. Two unstable modes are clearly visible for $ngh = 5$ and almost disappear for $ngh = 11$ (weak instabilities with a positive imaginary part of the order of 10^{-9} are present). In fact, the numerical wave should travel at a distance (in terms of the number of mesh points) no larger than the CFL number at each time step.

Additionally, the above-mentioned simplifications can introduce an error in the interface region. Such an

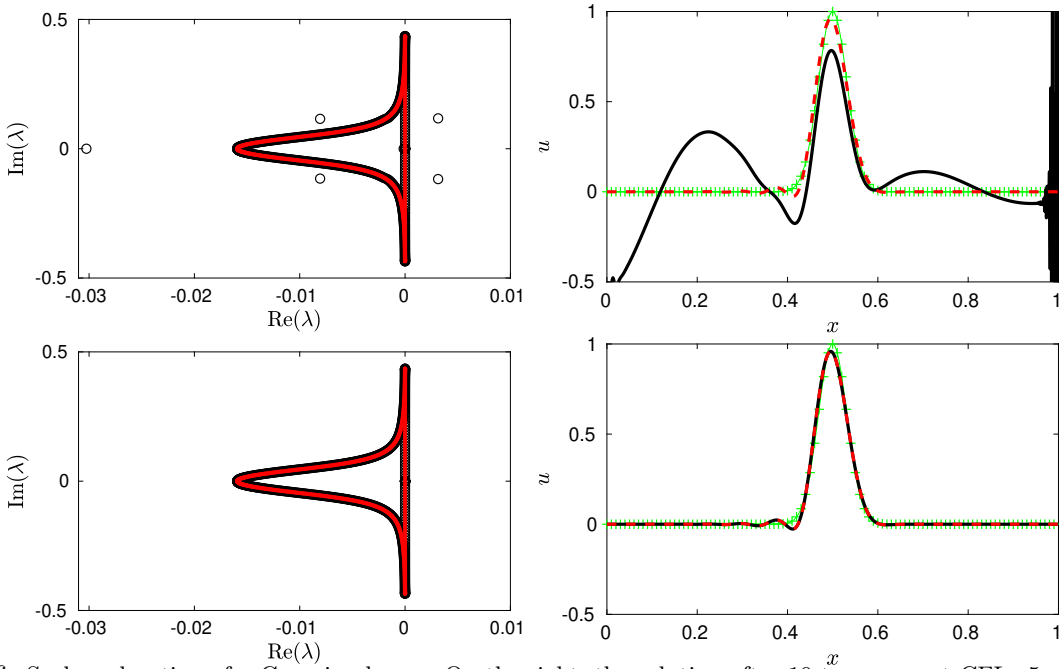

Figure 6: Scalar advection of a Gaussian hump: On the right, the solution after 10 turnovers at CFL=5 using true periodic conditions (– – –) and the approximate conditions with ghost points (——) are compared to the initial (exact) solution (++). The corresponding eigenspectrum is given on the left with red symbols for the periodic problem et black symbols for the ghost-point problem. For top figures, $ngh = 5$ and for bottom figures, $ngh = 11$.

error can be reduced by increasing the number of ghost cells from 5 on each side of the domain to a given integer ngh, so that a reasonable tradeoff between cost and accuracy has to be found. Previous analyses by Wu and Zou [35] for a time-lagging interface condition suggest that "since the CFL number is based on the maximum wave speed (eigenvalue), it is natural that a local perturbation (due to time-lagging) of the scheme will travel at a distance, in terms of the number of mesh points, no larger than the CFL number at each time step." An optimal overlapping width of 2CFL+1 would enable to contain the error in the overlapping region. As a consequence the overlapping method using ghost points should maintain the order of accuracy of the interior scheme if $ngh \approx$2CFL+1. This point will be investigated in § 5.1 for a vortex advection problem. In practice, all following applications are run on parallel computers. The maximum CFL is tested for each interface and the number of ghost points is set to max(5,2CFL+1). Generally, high CFL region are located near wall boundaries and few interfaces are concerned by the ghost-point extension. Each domain can have a variable ngh on each phase and the communication management is done with MPI one-sided RMA protocol.

5 Numerical applications

5.1 Vortex advection

As noted earlier, the major issue in multi-domain computations is the treatment of interfaces between domains. In the present work, communication between adjacent meshes is conducted through finite-size overlaps. Their effect on solution accuracy and stability is investigated initially for the unsteady inviscid flow generated by the a vortex convected by a uniform flow with freestream Mach number $M = U_\infty/c = 0.5$. The initial condition is determined from Taylor's theoretical vortex model:

$$\begin{cases} u = U_\infty + A\frac{y}{\Delta y}\exp(\alpha R^2) \\ v = -A\frac{x}{\Delta x}\exp(\alpha R^2) \\ p = p_\infty - \rho_\infty \frac{A^2}{4\alpha\Delta x\Delta y}\exp(2\alpha R^2) \end{cases} \tag{31}$$

where $R = \sqrt{(x-x_0)^2+(y-y_0)^2}$ with $(x_0, y_0) = (-30, 0)$ the initial vortex position, using a reference length $L = 1$ m. The vortex strength is $A = 5$ and $\alpha = -\ln 2/b^2$, where the Gaussian half-width is set to $b = 4$. The domain considered extends $-100 < x < 100$, $-100 < y < 100$ discretized by $N \times N$ uniformly spaced points, and periodic conditions are enforced in both directions. The grid is subdivided into four equally sized subdomains. The initial vortex location is located on the interface $y = 0$ and the vortex is advected from left to right during 100 time units L/U_∞. It crosses the interface located at $x = 0$.

The L_2-norm of the error with respect to the analytical solution is first used to assess the solver accuracy. The first series of runs, reported on the left of Fig.7, uses low CFL values between 0.02 and 1 with explicit time marching and deliberately coarse grids ranging between $N = 50$ to 300 to distinguish the accuracy of the spatial scheme. A slope of 2.5 in the log-log plot is obtained at CFL=1 and a very small time step has to be used to approach the spatial scheme accuracy. A slope of 7 is obtained for CFL=0.02. This first series shows that the temporal error rapidly dominates for this advection case.

In the second series of runs (middle plot of Fig.7), the implicit solver is used for large CFL numbers between 1 and 10 on fine grids ($N = 200$ to 1600) to highlight temporal integration errors. At CFL=1, the second-order accuracy of RK4 for nonlinear problems is recovered, and the slope slightly increases for higher CFL. The error saturates for the coarsest grids and CFL 8 to 10, since the error level is very high and the vortex is severely damaged during its advection.

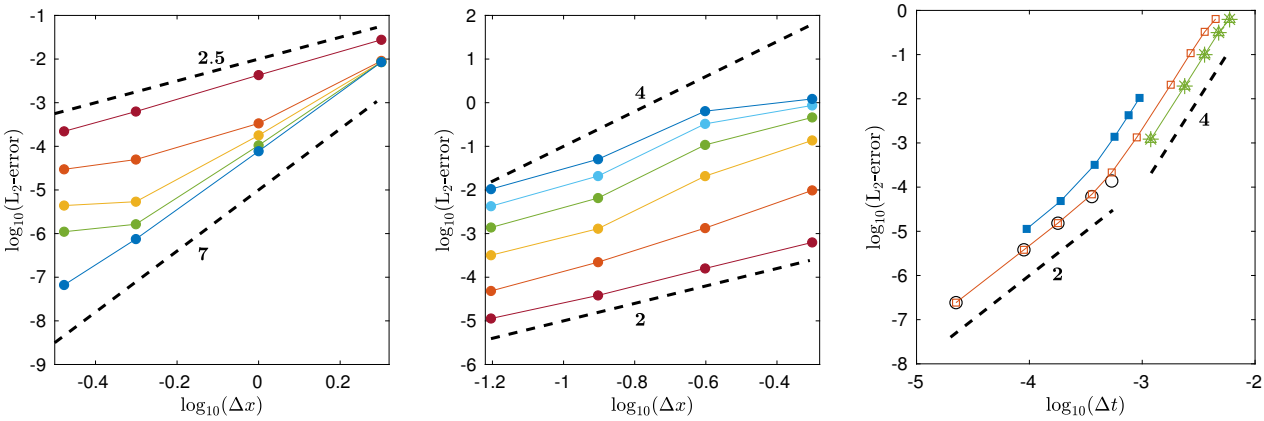

Figure 7: Vortex advection: L_2-norm of the error with respect to the analytical solution. Left, error of the explicit solver as a function of the mesh size Δx (grids 50×50, 100×100, 200×200, 300×300) for small CFL numbers (●,1; ●,0.5; ●,0.1;●,0.1; ●,0.02). Middle, error of the implicit solver as a function of the mesh size Δx (grids 200×200, 400×400, 800×800, 1600×1600) for large CFL numbers (●,1; ●,2; ●,4;●,6; ●,8; ●,10). Right, error as a function of the time step Δt: explicit solver (○) and IRS4 solver (□) for the grid 400×400 on 4 blocks at M=0.5; IRS4 solver on grid 400×400 at M=0.1 on 4 blocks (△) and 1 block (∗); IRS4 solver on grid 1600×1600 at M=0.5 (■).

Finally in the right subfigure of Fig.7, error logarithm is plotted as function of logarithm of Δt for the grid 400×400 and CFL ranging from 0.1 to 10 with IRS4 smoothing. Up to CFL 1.2 the error is the same as the one obtained with the explicit solver (black circles) and is second-order accurate. For higher values, a steeper slope of 4 is observed since the advection error rapidly rises. By reproducing the same numerical experiment with a very fine grid ($N = 1600$) for CFL 1, 2, 4, 6, 8, 10 (filled squares), it is clear that the change of slope is due to the difficulty to advect information with large time steps. A last test for the large CFL values is realized with $N = 400$ for a lower Mach number, $M = 0.1$. The error levels are the same as the ones obtained for $M = 0.5$. For these runs, a single domain was also used, which does not change the measured error, meaning that the errors are not related to the interface treatment.

In the preceding tests, the number of ghost cells for the MPI domain overlaps has been set to 2CFL+1. The influence of the number of ghost cells is illustrated in Fig.8 for $N = 1600$ and CFL=10. We report the iso-contours of the fluctuating pressure field, obtained by varying the number of ghost cells from 6 to 18. Intense spurious noise is generated when 6 ghost cells are used (note that the simulation blows up for 5 ghost cells). The spurious acoustic pulse and the vertical and horizontal oscillations developing at the interfaces are reduced for 9 ghost cells and are almost absent when 18 ghost cells are used (the theoretical value is 2CFL+1=21).

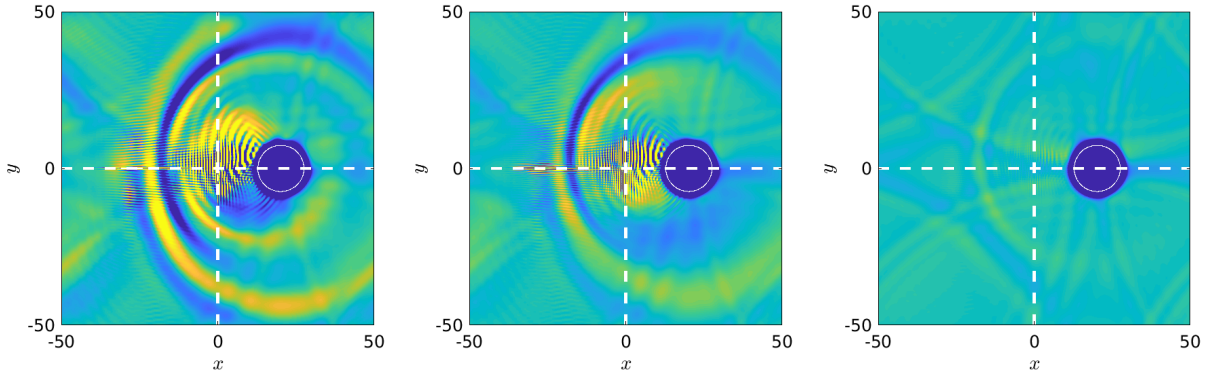

Figure 8: Vortex advection: maps of fluctuating pressure (between ±5 Pa, and 1 white contour for -100 Pa identifying the vortex core) on the grid 1600×1600 and CFL=10 using 6 (left), 9 (center) and 18 (right) ghost points. The white dashed lines mark the domain interfaces.

5.2 DNS of turbulent channel flow

The increase of CFL thanks to IRS4 smoothing is particularly beneficial for the simulation of wall-bounded turbulent flows, where very small mesh sizes are required near the wall to capture the stiff velocity gradient and predict the wall friction accurately. The first case investigated is the turbulent channel flow at Reynolds number $Re_\tau = (\rho_w u_\tau H)/\mu_w = 180$, based on the friction velocity u_τ, the channel half-height H and the wall density and viscosity ρ_w and μ_w, respectively. The computational domain of $4\pi H \times 2H \times 2\pi H$ is discretized with a grid $192 \times 180 \times 160$, uniformly spaced in the homogeneous directions x and z and clustered at the wall in the wall-normal direction y. This leads to a DNS resolution of $\Delta x^+ = 11.9$, $\Delta z^+ = 7.1$, $\Delta y_w^+ = 0.8$ and $\Delta y_c^+ = 4$, where the subscripts w and c are used to denote the near-wall and centerline resolutions, respectively. Isothermal no-slip conditions are applied at the walls and periodicity conditions along the streamwise and spanwise directions. The Mach number is set to $M = 0.3$ and the simulation is parallelized on 80 cores. For the implicit cases, IRS4 is applied in the wall-normal direction only, and the number of ghost points used for the IRS4 at the MPI interfaces is set to 5, i.e. the minimum required by the scheme stencil. Results using an explicit time integration with a global time step Δt_{expl} are compared with IRS4-accelerated simulations using global time steps ranging from $4\Delta t_{expl}$ to $7\Delta t_{expl}$ (see Table 1).

Case	Time integration	CFL_{max}	Δt^+	Nb iterations	$\frac{t_{CPU,expl.}}{t_{CPU,case}}$	$\frac{t_{CPU/it/proc,case}}{t_{CPU/it/proc,expl}}$	Legend
EXPL	RK4, explicit	1	1.05×10^{-2}	750 000	1.0	1.0	·····
IMPL4	RK4, IRS4	4	4.21×10^{-2}	187 000	3.30	1.18	- - -
IMPL5	RK4, IRS4	5	5.26×10^{-2}	150 000	4.06	1.18	- · - ·
IMPL6	RK4, IRS4	6	6.31×10^{-2}	125 000	5.08	1.18	- · - · ▽
IMPL7	RK4, IRS4	7	7.36×10^{-2}	110 000	5.93	1.18	- · - · ○

Table 1: Numerical parameters and computational performance of turbulent channel flow simulations with and without IRS4.

In the explicit case, the maximum CFL number, CFL_{max} is chosen equal to 1 to ensure stability throughout the simulation. This corresponds to a time step of $\Delta t^+ = 1.05 \times 10^{-2}$. Choi & Moin [37] studied the effects of the computational time step on the numerical solutions for an incompressible turbulent channel flow at $Re_\tau = 180$ using a fully implicit method. Up to $\Delta t^+ = 40 \times 10^{-2}$, no significant deterioration of turbulence statistics was observed. For the case IMPL7, which corresponds to the highest CFL_{max} used here, the time step is equal to 7.36×10^{-2}, well below the limit found in [37], so we expect the physical time step to have little impact on the accuracy of the solution. This allows to observe the influence of additional numerical errors introduced by the IRS4 on the solution quality.

For all implicit cases, the total number of iterations used to reach a statistically converged solution corresponds to the total number of iteration of case EXPL divided by CFL_{max}, so that the statistics are collected over the same physical time interval for all cases. Comparisons of the computational cost for the

explicit and implicit cases show that applying IRS4 in the wall-normal direction represents an overcost of 18% of CPU time (per point and per iteration). However, the total number of iterations being reduced proportionally to the CFL_{max} in use, the total computation time is reduced. For instance a reduction by a factor 4.06 is obtained for case IMPL5, and a factor 5.93 is found for IMPL7.

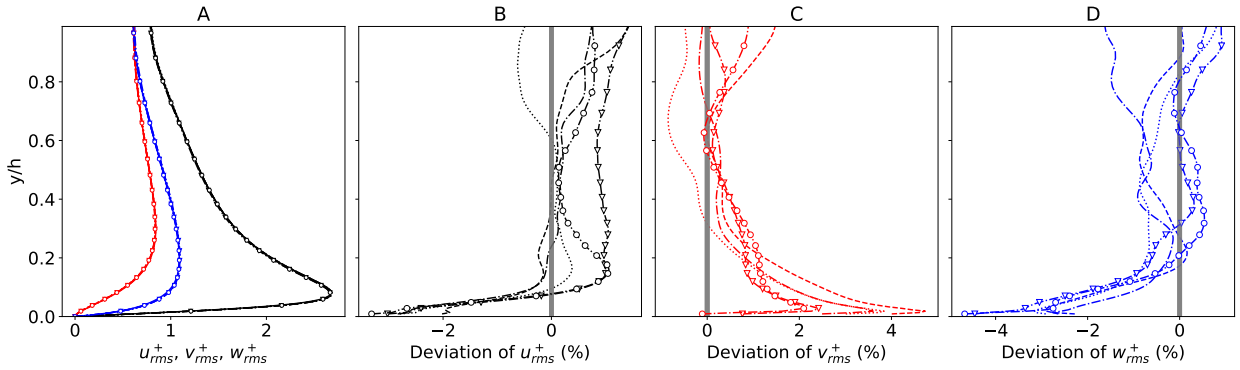

Figure 9: Comparison of the rms velocities profiles (u^+_{rms}, v^+_{rms} and w^+_{rms}) with the Vreman & Kuerten (——) reference [38] for the explicit and IRS4 time integration cases. Line legends in Table 1.

Figure 9 shows numerical solution for the various cases at stake. These are compared to the reference solution of Vreman & Kuerten [38], obtained with an incompressible solver based on spectral method and a very fine grid. The root-mean-square (rms) velocities profiles are plotted in Figure 9(A) along with the reference solution and we observe a very good agreement for all cases. To highlight differences, the relative deviations from the reference are plotted Fig.9(B, C and D). These are defined as $\mathrm{dev}(u^+_{i,rms}) = (u^+_{i,rms,\mathrm{VK}} - u^+_{i,rms})/u^+_{i,rms,\mathrm{VK}}$ with $u^+_{i,rms,\mathrm{VK}}$ the solution of Vreman & Kuerten for the i-th velocity component. Except for points close to the wall, deviations are below 1% for all cases. Moreover, the application of the IRS4 does not lead to any noticeable increase of the relative deviation. For an accumulation of statistics on an equivalent duration, we observe a good convergence of second-order statistics with a reduced computational cost when applying IRS4. For simulations that require a longer time to converge statistics or to obtain higher-order statistics, introducing IRS4 acceleration is therefore particularly beneficial.

5.3 LES of real-gas turbulent boundary layers

As a second application we consider boundary-layer transition of a real gas at high-subsonic condition (Mach 0.9). For this spatially evolving configuration, time-step constraints for DNS resolution of the wall-bounded turbulence as similar to those encountered in the channel flow case. In the laminar and transitional regimes, growing instability waves are particularly sensitive to dispersion errors, whereas numerical resolution of the turbulent region strongly depends on numerical dissipation. A modal transition is performed by injecting at the flow inlet plane a pair of oblique modes skewed by 30° with respect to the streamwise direction and with a nondimensional angular frequency $\omega_0 L^*/U_\infty = 0.02$. The inlet Reynolds number Re_{L^*} based on Blasius length $L^* = x/\sqrt{Re_x}$ is equal to 1000. The spanwise extent is taken equal to one spanwise wavelength of the input modes, $\lambda_z = 2\pi/\beta_0$, with the spanwise wavenumber $\beta_0 = 0.04L^*$. Adiabatic no-slip conditions are applied at the wall, and non-reflecting Tam & Dong's conditions are imposed at the inlet, top and outflow boundaries. The working fluid is perfluorinated ketone Novec$^{\mathrm{TM}}$649, a dense gas used in energy conversion systems, and specifically Organic Rankine Cycles (ORC). The thermodynamic flow conditions correspond to a freestream temperature of 100°C and a pressure of 4 bars. At such conditions, the fluid deviates from ideal gas behavior. Real-gas behavior is modeled with the Peng-Robinson-Stryjek-Vera equation of state [39] and the Chung-Lee-Starling model [40] for the transport properties.

First, a DNS with the explicit time integration is realised over a computational grid of $9000 \times 400 \times 1000$ points, achieving a total of 3.6 billions points parallelized on 16 384 processors. The resolution in wall units corresponds to Δx^+=14 in the streamwise direction, Δz^+=7 in the spanwise direction, and Δy^+_w=0.7 and Δy^+_e=11 at the wall and at the boundary-layer edge, respectively. The grid is uniform in the streamwise

[t]

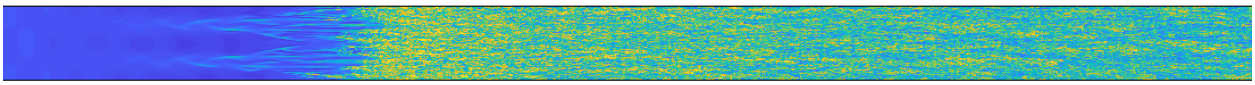

Figure 10: Instanteaneous view of the streamwise velocity at $y^+ \approx 15$ for the DNS of real-gas boundary-layer flow.

direction and the spanwise direction, and streched in the direction normal to the wall by a factor of 1.015 up to the 190^{th} point. The streamwise domain extends up to a momentum-thickness Reynolds number $Re_{\theta,end} \sim 5000$, allowing the observation of high Reynolds-number effects. Despite the high subsonic Mach number, we found that the mean and fluctuating turbulent profiles are very close to incompressible DNS databases, essentially due to the high specific heat of Novec$^{\text{TM}}$649 which significantly reduces friction heating at the wall.

Case	CFL	$\Delta t/\Delta t_{\text{DNS}}$	$N_x \times N_y \times N_z$	Δx^+	Δy_w^+	Δy_e^+	Δz^+	$Re_{\theta,end}$	Legend
DNS	1	1.00	$9000 \times 400 \times 1000$	14	~ 0.7	~ 11	7	~ 5000	—
LES-expl	1	1.33	$3000 \times 360 \times 660$	28	~ 0.8	~ 14	11	~ 3300	- - -
LES-IRS	4	5.33	$3000 \times 360 \times 660$	28	~ 0.8	~ 14	11	~ 3300	-·-·-

Table 2: Characteristics of the different cases for the real-gas turbulent boundary layers

Second, two LES are carried out at the same conditions as the DNS, in order to validate our LES strategy. The LES computational grid, composed by 712 million points, is obtained by reducing the resolution and the streamwise length (see Table 2). We chose a wall-resolved implicit large eddy simulation strategy, whereby the explicit selected filter ensures selective regularization of the subfilter turbulent scales. This implicit modeling strategy has been shown to be effective [16] and avoids the computational overhead introduced by the explicit subgrid-scale models. The first LES, using explicit time marching at CFL_{max}=1 is referred to as LES-expl hereafter and is used to validate our wall-resolved LES strategy, in particular the resolution and the implicit subgrid model. For the second LES, called LES-IRS, the IRS4 smoothing is applied in the wall-normal direction with a CFL_{max} of 4, resulting in computational time reduction by a factor 3.5 with respect to the explicit case. The main characteristics of the different simulations are summarized in Table 2. A snapshot of the boundary-layer transition is reported in Fig.10. For more details, the reader is referred to [41].

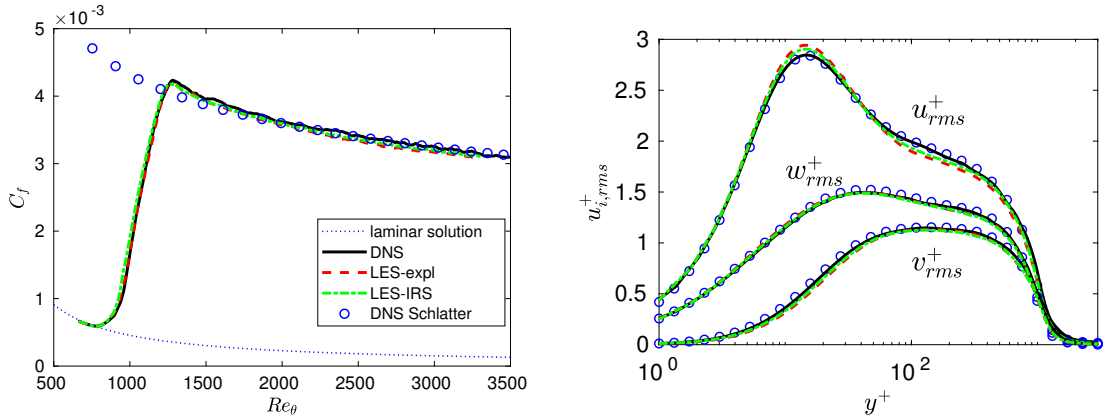

Figure 11: Skin friction coefficient (left) and rms velocities profile at Re_θ=3270 (right) for the different cases and the DNS of Schlatter and Örlü [42].

Since the turbulent boundary layers exhibits an incompressible behavior due to the high specific heat of the real gas, the results are compared with the incompressible DNS of Schlatter and Örlü [42]. Results for the friction coefficient C_f are reported in Fig.11. The C_f evolution of both LES calculations are in very good

agreement with the DNS, including the location of the transition. We focus in the following on the turbulent boundary layer at the location Re_θ=3270. The turbulent intensity profiles are reported in Fig.11. A good match with the DNS and the incompressible reference is observed, with a slight underestimation of turbulent intensities, due to the LES resolution. The profiles for LES-expl and LES-IRS are almost superimposed, demonstrating the accuracy of the implicit time advancement.

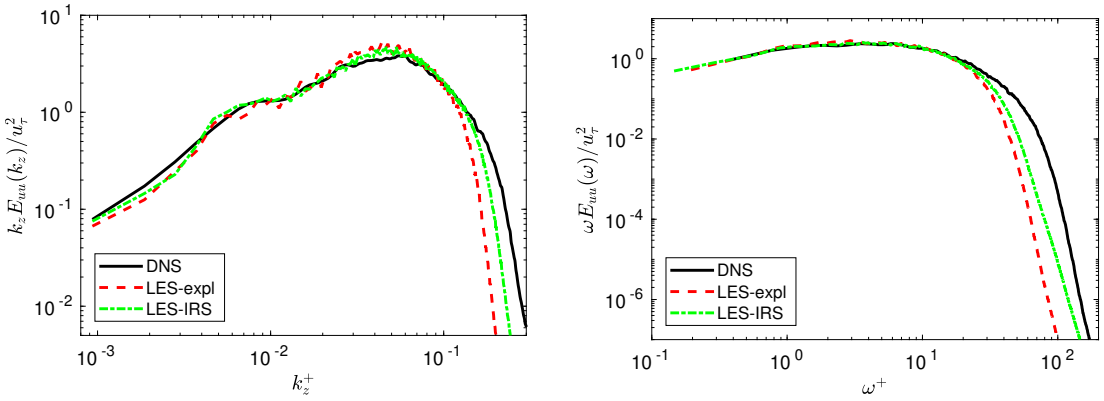

Figure 12: Premultiplied spanwise (left) and temporal (right) spectra of the streamwise velocity fluctuations for the turbulent boundary layer at Re_θ=3270.

To investigate in more detail the differences between the cases, the spanwise and temporal spectra of the streamwise velocity fluctuations are plotted in Fig. 12. As expected, the premultiplied spanwise spectra exhibit an earlier cut-off for the two LES, due to coarser spanwise resolution. Interestingly, we observe a better resolution of LES-IRS w.r.t LES-expl at the small scales. This is a consequence of the lower numerical dissipation introduced in LES-IRS, because numerical dissipation is applied less frequently over the integration interval, due to the larger time step. Similar results are obtained for the temporal spectra. This means that the LES cut-off is governed by streamwise resolution and not by time resolution. Finally, we observe a smoother decaying slope for LES-IRS, most likely due to the dispersion of the IRS operator.

5.4 Laminar flow past a cylinder

We now investigate the performance of IRS4 on curvilinear grids. The first test problem is unsteady laminar flow around a 2D circular cylinder at a Mach number M=0.3 and a Reynolds number Re_D=1200, based on cylinder diameter D and freestream velocity U_0. Even if the flow is expected to become three-dimensional at the selected Reynolds number, this configuration remains a good numerical test case used in the literature to assess the performance of numerical schemes [43, 44]. In the aim of observing the influence of time integration errors, a relatively fine grid of 360×300 points is selected. An O-grid is generated with a first mesh size at the wall of $0.0028D$ and free boundary located at a distance of $20D$. Grid points are uniformly distributed in the azimuthal direction. Tam & Dong's non reflecting conditions are used at the free boundary and an adiabatic wall condition is prescribed. The initial flow is symmetric and the transient phase is run with the explicit solver (CFL_{max}=0.5) during 600 000 iterations. The recirculation bubble behind the cylinder starts to grow, then becomes asymmetric and a regular von Kármán vortex shedding in the wake is finally established. Starting from the same established flow solution, 7 simulations are performed over the same physical time interval using increasing maximum CFL numbers: 200 000 additional iterations are run with the explicit solver at CFL_{max}=0.5; then, plugging the IRS4 smoothing, 100 000 at CFL_{max}=1, 50 000 at CFL_{max}=2, 25 000 at CFL_{max}=4, 12 500 at CFL_{max}=8, 8 333 at CFL_{max}=12 and 6 250 at CFL_{max}=16.

Close-up views of the vorticity in the cylinder wake are plotted in Fig.13 for the two extreme CFL numbers (intermediate cases are not reported for brevity). The snapshot is taken after one quarter of the number of iterations. The explicit case will serve as reference in the following. At first glance, the vorticity field is little affected even for a time step multiplied by 32. In fact, we show below that the dispersive error of IRS essentially affect the shedding frequency, and the vortices are slightly shifted. The lift coefficient history is presented in Fig.14, and the progressive frequency shift is visible for the two highest CFL values. Looking

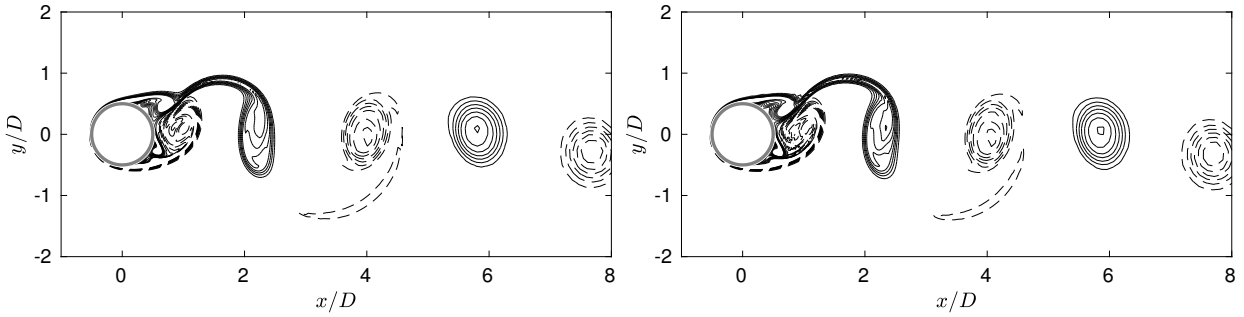

Figure 13: Laminar flow past a cylinder at $Re_D = 1200$. Snapshot of the vorticity field at the same instant for the explicit at CFL=0.5 (left) and implicit IRS4 at CFL=16 (right). 8 positive (———) and negative (– – –) isocontours of $\omega_z D/U_0$ from ± 1 to 8.

more closely at the last shedding cycles for lift and drag coefficients, we can see that the accumulation of phase errors is only significant for CFL_{max} 12 and 16. The quantitative errors with respect to the low-CFL explicit case are reported in Table 3. The deviations for the Strouhal number fD/U_0 are not measurable for CFL_{max}=1 and 2 because the power spectral density is performed for 13 cycles. An error of more than 10% is noted for CFL=16. The errors for mean and fluctuating aerodynamic coefficients remains low below 2% for the lift and below 0.3% for the drag. The logarithms of the solution errors in the lift, drag mean and rms values and Strouhal number are shown in Fig.15 as a function of logarithm of time step. We can see that the error grows at a greater rate for high CFL values. The slope modification occurs for CFL around 10. Finally, an idea of the reduction in computational cost by inspection of the right panel of Fig.15. For the present curvilinear setting, the IRS overcost in 2D is roughly of 50% per implicited direction, leading to a total overcost of 100% (the cost of the explicit run at CFL_{max}=0.5 is approximately the same as the implicit run at CFL_{max}=1). We can see that the cost decrease is linear in a log-log plot and that a saving of a factor of 5 can be achieved with a solution of comparable accuracy with our reference.

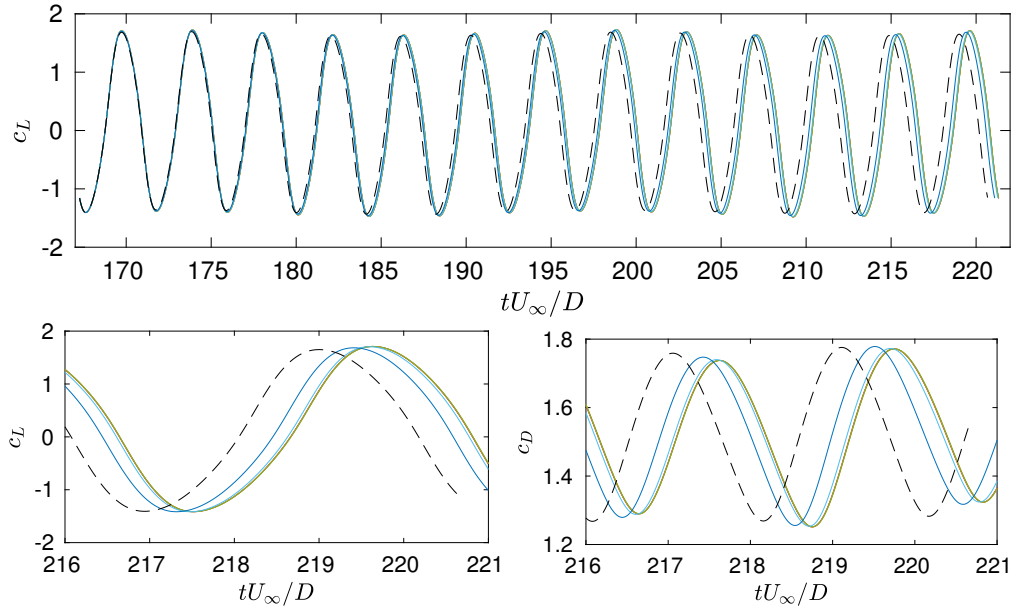

Figure 14: Laminar flow past a cylinder at $Re_D = 1200$. Top, lift evolution for various CFL numbers (———,0.5; ———,1; ———,2; ———,4; ———,8; ———,12; – – –,16). Bottom, close-up views for the lift and drag for the last simulated cycles.

CFL	0.5	1	2	4	8	12	16
$c_{L_{\text{rms}}}$	1.1231	1.1230	1.1230	1.1229	1.1215	1.1153	1.1030
(error)	(0%)	(-0.0083%)	(-0.0125%)	(-0.0176%)	(-0.1417%)	(-0.6987%)	(-1.7901%)
$c_{D_{\text{mean}}}$	1.5327	1.5328	1.5329	1.5332	1.5336	1.5323	1.5282
(error)	(0%)	(+0.0021%)	(+0.0088%)	(+0.0309%)	(+0.0565%)	(-0.0288%)	(-0.2936%)
$c'_{D_{\text{rms}}}$	0.1674	0.1674	0.1673	0.1673	0.1673	0.1675	0.1677
(error)	(0%)	(-0.0151%)	(-0.0192%)	(-0.0361%)	(-0.0492%)	(+0.0732%)	(+0.2142%)
St	0.224	0.224	0.224	0.230	0.233	0.238	0.248
(error)	(0%)	(0%)	(0%)	(+2.5641%)	(+3.8961%)	(+6.3939%)	(+10.7266%)

Table 3: Errors with respect to the explicit solution for the laminar flow past a cylinder at $Re_D = 1200$.

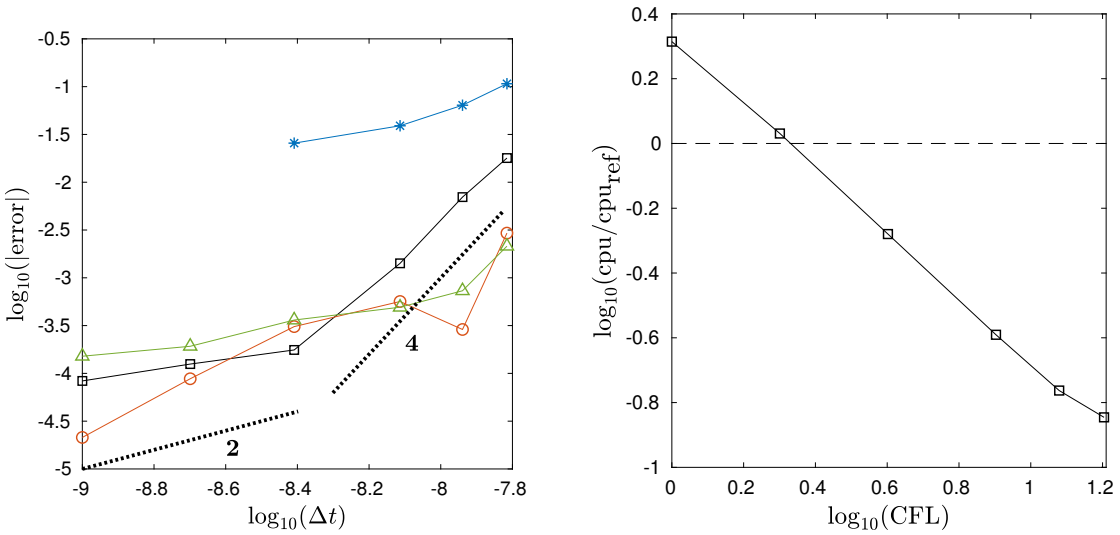

Figure 15: Laminar flow past a cylinder at $Re_D = 1200$. Left, error for increasing CFL numbers for the lift fluctuations (——□), mean drag (——○), drag fluctuations (——△) and Strouhal number (——∗). Right, CPU cost compared to explicit time integration.

5.5 LES of turbulent flow past a cylinder

The IRS4 method is finally validated for the flow past a circular cylinder at $Re_D = 3900$ based on the diameter D, and at $M = 0.3$, which is a common benchmark case for curvilinear geometries [45, 46]. The simulation is performed on a multi-block H-O-H grid topology with approximatively 10 millions points parallelized on 102 processors. 270 points are used around the cylinder with a first mesh size of $0.002D$, 96 points are used in the spanwise direction to discretize $2D$ and the wake extent is discretized by 405 points. Non-reflecting Tam & Dong's conditions are applied at free boundaries and a sponge zone is added at the outlet boundary. An explicit simulation is started from the initial field described in the AS1 benchmark case during 500 000 iterations with a nondimensional time step $\Delta t U_0/D = 4.55 \times 10^{-4}$. The explicit calculation is then run for 500 000 and statistics are accumulated. A snapshot of the instantaneous field is given in Fig.16.

Afterwards, three implicit simulations with IRS4 are run by multiplying the time step by 2,3,4 and dividing the number of iterations by the same amount. The IRS4 smoothing is applied in the ξ and η-directions, yielding an overcost of approximately 20% per direction. A saving in CPU time of a factor of 1.44, 2.16 and 2.98 is obtained. Time-averaged fields, computed over the same physical duration, are in good agreement, as shown in Fig.17 and 18. The explicit and the implicit case with $4\Delta t$ are almost superimposed and in very good agreement with the experiment of Parnaudeau *et al.* [45].

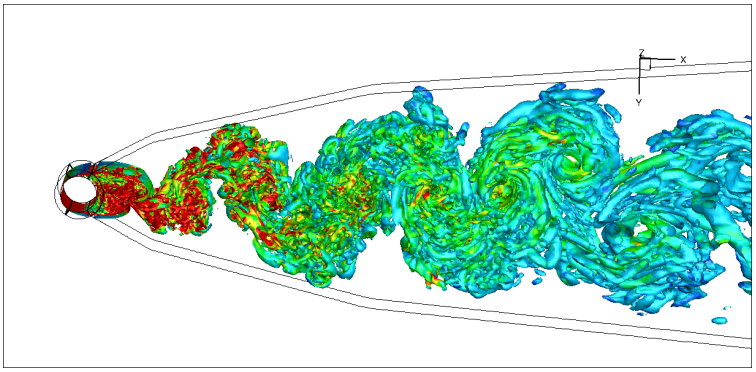

Figure 16: Flow past a cylinder at $Re_D = 3900$. Snapshot of λ_2 vorticity criterion.

6 Conclusion and Future Work

Fourth-order implicit residual smoothing (IRS4) is used for time implicitation of the compressible Navier-Stokes on multi-block curvilinear meshes. The method is applied to DNS and LES of turbulent flows and high-accuracy finite-difference schemes are used in conjunction with a suitable numerical dissipation. IRS4 acts as a filter on the solution increments at each substep of the underlying Runge–Kutta algorithm. It changes the spectral properties of the spatial operator and essentially introduces a dispersive error by contracting the locus of the spatial operator near the imaginary axis. In order to enhance robustness at high CFL in the context of high-order schemes, the main modifications of the method with respect to the original version in Ref.[1] are:

- a filtering strategy is used instead of the artificial viscosity, which was part of the spatial scheme and was thus affected at high CFL values by the strong contraction that reduces severely the effective dissipation near the grid cut-off. On the contrary, filtering is independent on the time advancement;
- the extension to curvilinear grids takes into account the coupling of mesh directions with the coordinate transform by using a spectral radius mixing both directions, which enters the IRS coefficients;
- the treatment of boundary conditions uses reduced-stencil IRS operators to fill the IRS matrix. Coefficients on the first row of the pentadiagonal matrix correspond to one-sided IRS1 and IRS2 coefficients are used for the second row;
- the parallel and multi-block implementation rely on an overlap between domain where ghost points are communicated. An optimal width of the overlap is found to be of the order of 2CFL+1.

Some simple to massively parallel applications are shown that demonstrate the efficiency of IRS4. A saving in terms of computational time of a factor 3 to 5 in 3D is expected for scale-resolving turbulent simulations. The overcost of IRS4 smoothing is approximately 15 to 30% per implicited direction. The greater time steps can also allow a better convergence of statistics, notably for low-frequency phenomena such shock wave-boundary layer interation. We are also currently applying the scheme to the compressible flow inside a fluidic actuator, which implies a relatively complex multi-block arrangement and requires the addition of shock-capturing since the flow is highly unsteady and supersonique. A future work will be to further assess the combination of filtering and low-order numerical dissipation to simulate flows with shocks.

Acknowledments

This research has been funded by the Agence Nationale de la Recherche through the ANR-20-CE92-0019-02 project Regal-ORC. This work was granted access to the HPC resources of IDRIS and TGCC under the allocation A0112A01736 made by GENCI (Grand Equipement National de Calcul Intensif).

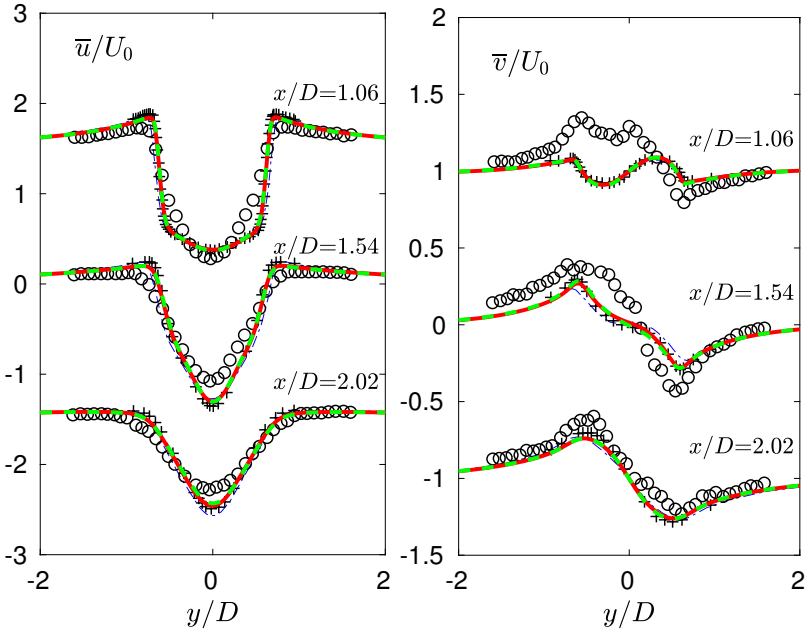

Figure 17: Flow past a cylinder at $Re_D = 3900$. Mean wake velocity profiles for the explicit (——) and implicit IRS4 (– – –) LES, compared to experiment (++) and LES (– · – ·) of Parnaudeau *et al.* [45], and experiment of Lourenco & Shih [47] (○○).

References

[1] P. Cinnella and C. Content. High-order implicit residual smoothing time scheme for direct and large eddy simulations of compressible flows. *J. Comput. Phys.*, 277:72–100, 2014.

[2] J.-C. Hoarau, P. Cinnella, and X. Gloerfelt. Large eddy simulation of turbomachinery flows using a high-order implicit residual smoothing scheme. *Computers and Fluids*, 198:104395–1–104395–13, 2020.

[3] A. Lerat. Une classe de schémas aux différences implicites pour les systèmes hyperboliques de lois de conservation. *C. R. Acad. Sci.*, 288A:1033, 1979.

[4] A. Lerat, J. Sidès, and V. Daru. An implicit finite-volume method for solving the Euler equations. In *Lecture Notes in Physics*, volume 170, pages 343–349, 1982.

[5] A. Lerat. Implicit methods of second order accuracy for the euler equations. *AIAA Journal*, 23(1):33–40, 1985.

[6] K. Khalfallah, G. Lacombe, and A. Lerat. Analysis of implicit treatments for centred Euler solver. *Computers and Fluids*, 22(2/3):381–406, 1993.

[7] A. Jameson and T.J. Baker. Solution of the Euler equations for complex configurations. *AIAA Paper 83-1929*, 1983.

[8] L. Martinelli and A. Jameson. Validation of a multigrid method for the Reynolds averaged equations. *AIAA Paper 88-0414*, 1988.

[9] B. van Leer, C.-H. Tai, and K.G. Powell. Design of optimally smoothing multi-stage schemes for the Euler equations. *AIAA Paper 89-1933*, 1989.

[10] P.C.E. Jorgenson and R.V. Chima. An unconditionally stable Runge-Kutta method for unsteady flows. *AIAA Paper 89-0205*, 1989. also NASA Technical Memorandum 101347.

[11] E. Turkel, R.C. Swanson, V. Vatsa, and J.A. White. Multigrid for hypersonic viscous two- and three-dimensional flows. *AIAA Paper 91-1572*, 1991. also NASA Contractor Report 187603, ICASE Report No.91-57.

[12] J. Blazek, N. Kroll, R. Radespiel, and C.-C. Rossow. Upwind implicit residual smoothing method for multi-stage schemes. *AIAA paper 91-1533*, 1991.

[13] J. Blazek. *Computational Fluid Dynamics: Principles and Applications*, chapter 9 - Acceleration Techniques, pages 283–335. Butterworth-Heinemann, Oxford, third edition edition, 2015.

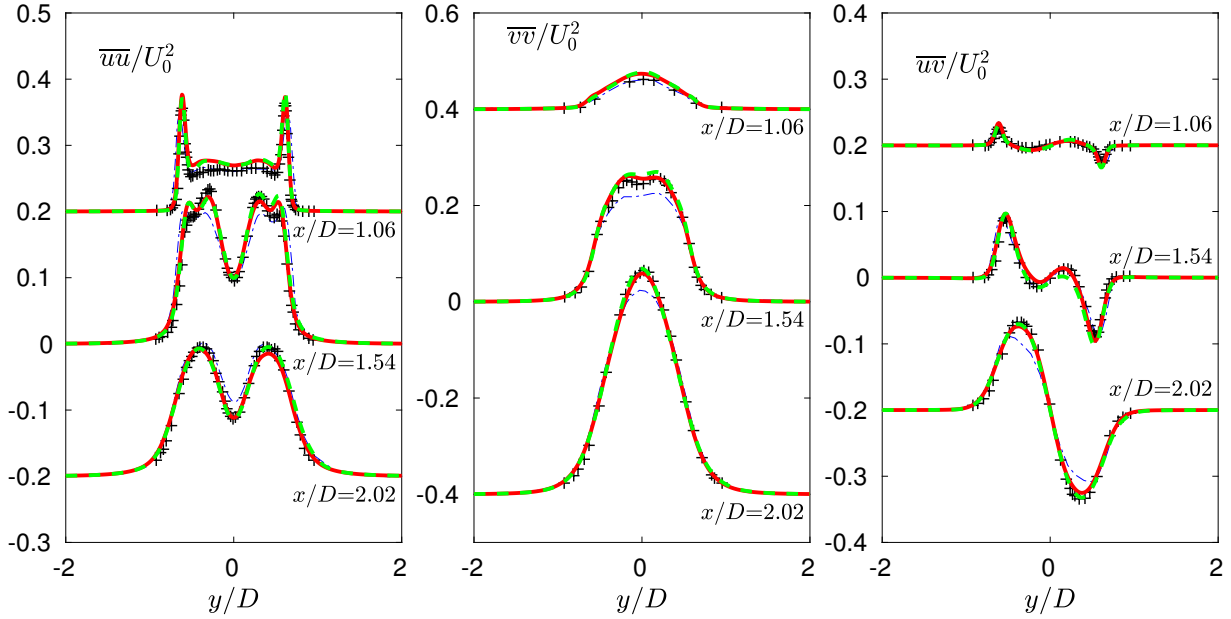

Figure 18: Flow past a cylinder at $Re_D = 3900$. Turbulent intensities profiles in the wake of the cylinder at $x/D = 1.06$, 1.54 and 2.02 for the explicit (——) and implicit IRS4 (– – –) LES, compared to experiment (++) and LES (– · – · –) of Parnaudeau *et al.* [45].

[14] A. Jameson, W. Schmidt, and E. Turkel. Numerical solutions of the Euler equations by finite volume methods using Runge-Kutta time-stepping schemes. *AIAA Paper 81-1259*, 1981.

[15] G. Aubard, P. Stefanin Volpiani, X. Gloerfelt, and J.-C. Robinet. Comparison of subgrid-scale viscosity models and selective filtering strategy for large-eddy simulations. *Flow, Turbulence and Combustion*, 91(3):497–518, 2013.

[16] X. Gloerfelt and P. Cinnella. Large eddy simulation requirements for the flow over periodic hills. *Flow, Turbulence and Combustion*, 103(1):55–91, 2019.

[17] S.K. Lele. Compact finite difference schemes with spectral-like resolution. *J. Comput. Phys.*, 103:16–42, 1992.

[18] D.V. Gaitonde and M.R. Visbal. Padé-type higher-order boundary filters for the Navier-Stokes equations. *AIAA Journal*, 38(11):2103–2112, 2000.

[19] A.K. Edoh, N.L. Mundis, C.L. Merkle, A.R. Karagozian, and V. Sankaran. Comparison of artificial-dissipation and solution-filtering stabilization schemes for time-accurate simulations. *J. Comput. Phys.*, 375:1424–1450, 2018.

[20] J.W. Kim and D.J. Lee. Adaptive nonlinear artifical dissipation model for computational aeroacoustics. *AIAA Journal*, 39(5):810–818, 2001.

[21] H.C. Yee, N.D. Sandham, and M.J. Djomehri. Low-dissipative high-order schock-capturing methods using characteristic-based filters. *J. Comput. Phys.*, 150:199–238, 1999.

[22] M. Ni, G. Xi, and S. Wang. Construction of high-order accuracy implicit residual smoothing schemes. *Applied Mathematics and Mechanics*, 21(4):407–414, 2000.

[23] P. Cinnella and A. Lerat. A fully implicit third-order scheme in time and space for compressible turbulent unsteady flow simulations. In *European Congress on Computational Methods in Applied Sciences and Engineering (ECCOMAS)*, Barcelona, Spain, 11-14 September 2000.

[24] T.H. Pulliam. Artificial dissipation models for the Euler equations. *AIAA Journal*, 24(12):1931–1940, 1986.

[25] J.C. Strikwerda. Initial boundary value problems for the method of lines. *J. Comput. Phys.*, 34:94–107, 1980.

[26] V. Couailler and N. Liamis. Unsteady Euler and Navier-Stokes flows simulations with an implicit Runge-Kutta method. In *Computational Fluid Dynamics'94*, pages 917–924. J. Wiley, 1994.

[27] B. Gustafsson. On the implementation of boundary conditions for the method of lines. *BIT Numer. Math.*, 38:293–314, 1998.
[28] M.G. Gasparo and S. Pieraccini. Implicit residual smoothing in a parallel 2D Euler solver. *Intern. J. Computer Math.*, 72:313–324, 1999.
[29] A. Povitsky and P.J. Morris. A high-order compact method in space and time based on parallel implementation of the Thomas algorithm. *J. Comput. Phys.*, 161:182–203, 2000.
[30] J. Choi and J.J. Dongarra. Scalapack: a scalable linear algebra library for distributed memory concurrent computers. *Proceedings of the Fourth Symposiumon the Frontiers of Massively Parallel Computation*, x:120–127, 1992.
[31] E. Polizzi and A. Sameh. SPIKE: A parallel environment for solving banded linear systems. *Computers and Fluids*, 36:113–120, 2007.
[32] K. Ki-Ha, K. Ji-Hoon, P. Xiaomin, and C. Jung-Il. PaScaL TDMA: A library of parallel and scalable solvers for massive tridiagonal systems. *Computer Physics Communications*, 260:107722, 2021.
[33] C. Borel and F.X. Roux. Implicit multi-domain method for unsteady compressible inviscid fluid flows around 3D projectiles. In K.G. Reinsch, W. Schmidt, A. Ecer, J. Häuser, and J. Periaux, editors, *Conference on Parallel Computational Fluid Dynamics,* 10-12 June 1991, pages 47–58, Stuttgart, Germany, 1992. Elsevier Science Publishers B.V.
[34] A. Lerat and Z.-N. Wu. Stable conservative multidomain treatments for implicit Euler solvers. *J. Comput. Phys.*, 123:45–64, 1996.
[35] Z.-N. Wu and H. Zou. Grid overlapping for implicit parallel computation of compressible flows. *J. Comput. Phys.*, 157:2–43, 2000.
[36] M.R. Visbal and D.V. Gaitonde. High-order-accurate methods for complex unsteady subsonic flows. *AIAA Journal*, 37(10):1231–1239, 1999.
[37] H. Choi and P. Moin. Effect of the computational time-step on numerical solutions of turbulent flow. *J. Comput. Phys.*, 113:1–4, 1994.
[38] A.W. Vreman and J.G.M. Kuerten. Comparison of direct numerical simulation databases of turbulent channel flow at re_τ=180. *Phys. Fluids*, 26:015102, 2014.
[39] R. Stryjek and J.H. Vera. PRSV: An improved Peng-Robinson equation of state for pure compounds and mixtures. *The Canadian Journal of Chemical Engineering*, 64(2):323–333, 1986.
[40] T.H. Chung, M. Ajlan, L.L. Lee, and K.E. Starling. Generalized multiparameter correlation for nonpolar and polar fluid transport properties. *Industrial & Engineering Chemistry Research*, 27(4):671–679, 1988.
[41] A. Bienner, X. Gloerfelt, and P. Cinnella. Numerical study of boundary-layer transition in a high-subsonic organic vapor flow. In *56th 3AF International Conference AERO2022*, Toulouse, France, March 28-30 2022. AAAF.
[42] P. Schlatter and R. Örlü. Assessment of direct numerical simulation data of turbulent boundary layers. *J. Fluid Mech.*, 659:116–126, 2010.
[43] H. Bijl, M.H. Carpenter, V.N. Vatsa, and C.A. Kennedy. Implicit time integration schemes for the unsteady compressible Navier-Stokes equations: Laminar flow. *J. Comput. Phys.*, 179:313–329, 2002.
[44] K. Grimich, B. Michel, P. Cinnella, and A. Lerat. An accurate finite-volume formulation of a residual-based compact scheme for unsteady compressible flows. *Computers and Fluids*, 92:93–112, 2014.
[45] P. Parnaudeau, J. Carlier, D. Heitz, and E. Lamballais. Experimental and numerical studies of the flow over a circular cylinder at Reynolds number 3900. *Phys. Fluids*, 20:085101, 2008.
[46] 4th International Workshop on High-Order CFD Methods, sponsored by NASA, AIAA, DLR and Army Research Office (ARO), June 3-4 2015 at the Eccomas / 6th European Conference on CFD (ECFD VI), Crete. Case AS1 `https://how4.cenaero.be/content/as1-dnsles-infinite-cylinder-re3900`
[47] L. M. Lourenco and C. Shih. Characteristics of the plane turbulent near wake of a circular cylinder, a particle image velocimetry study. Published in Ref.[48], 1994.
[48] P. Beaudan and P. Moin. Numerical experiments on the flow past a circular cylinder at sub-critical reynolds number. CTR Annual Research Briefs, NASA Ames/Stanford University, 1994.