Recycling Krylov Subspaces for Efficient Partitioned Solution of Aerostructural Coupled Adjoint Systems

Mehdi Jadoui, Christophe Blondeau, François-Xavier Roux

Recycling Krylov Subspaces for Efficient Partitioned Solution of Aerostructural Adjoint Systems.

Mehdi Jadoui* and Christophe Blondeau†
ONERA, Université Paris-Saclay, F-92322 Châtillon, France

François-Xavier Roux‡
Sorbonne Université, F-75005 Paris, France

Robust and efficient solvers for coupled-adjoint linear systems are crucial to successful aerostructural optimization. Monolithic and partitioned strategies can be applied. The monolithic approach is expected to offer better robustness and efficiency for strong fluid-structure interactions. However, it requires a high implementation cost and convergence may depend on appropriate scaling and initialization strategies. On the other hand, the modularity of the partitioned method enables a straightforward implementation while its convergence may require relaxation. In addition, a partitioned solver leads to a higher number of iterations to get the same level of convergence as the monolithic one. The objective of this paper is to accelerate the partitioned solver by considering techniques borrowed from Krylov subspace recycling strategies adapted to sequences of linear systems with varying right-hand sides. Indeed, in a partitioned framework, the structural source term attached to the fluid block of equations affects the right-hand side with the nice property of quickly converging to a constant value. We will also consider error approximation and approximate eigenvectors deflation in conjunction with advanced inner-outer Krylov solvers for the fluid block equations. We demonstrate the benefit of these techniques by computing the coupled derivatives for an aeroelastic configuration of the ONERA-M6 fixed wing in transonic flow. For this exercise the fluid grid was coupled to a structural model specifically designed to exhibit a high flexibility. All computations are performed using RANS flow modeling and a fully linearized one-equation Spalart-Allmaras turbulence model.

I. Introduction

NOWADAYS, aeronautical industries are deeply concerned about the environmental issues generated by the growth in air traffic. As such, they are constantly confronted with new technological challenges, either by improving the performance of an aircraft, or by making new insights. One of them consists in finding, in the early design stage, the best potential design configuration that minimizes a certain objective function (e.g the lift-to-drag ratio) with respect to a set of design variables leading potentially in the end to a great reduction of CO_2 emission. Fine control of an aerodynamic shape or of a structural layout leads to a high-dimensional parameter space. For high-fidelity simulations, gradient-based optimizers in conjunction with the adjoint approach are the methods of choice. However, the coupled adjoint linear system is inherently ill-conditioned as it embeds matrix blocks of different scales and structures. In addition, the fluid block coming from the exact linearization of the RANS equations associated to a turbulence model is often very stiff. Besides, a strong level of fluid-structure interaction is known to be detrimental to the robustness and efficiency of existing solution techniques. Resolving such linear systems is achieved by two main approaches. The partitioned (or segregated) approach and the monolithic approach. The first one simply consists in resolving, in an alternating way, the aerodynamic and the structural sub-problems by applying the Linear Block Gauss-Seidel algorithm (LBGS). It expresses the inder-disciplinary coupling as a source term to the right-hand side of each set of disciplinary adjoint equations. The modularity of this approach makes it rather interesting since it takes advantage of the specific routines designed for each sub-problem and does not demand a high implementation cost. Nevertheless, this approach becomes rapidly inefficient and could even diverge for strong fluid-structure coupling even though the addition of some

*PhD student, Aerodynamics, Aeroelasticity and Acoustics department, mehdi.jadoui@onera.fr.
†Research Scientist, Aerodynamics, Aeroelasticity and Acoustics department, christophe.blondeau@onera.fr
‡Professor, Laboratory Jacques-Louis Lions, roux@ann.jussieu.fr

degree of relaxation helps to mitigate this issue. On the other hand, the monolithic approach consists in solving the fluid and structural equations simultaneously making it more robust in the sense that it is less sensitive to the strength of the fluid-structure interaction. The coupled adjoint system is generally solved by using Krylov subspace methods. The challenging aspect of such an approach is then to develop advanced preconditioning strategies combined with numerical ingredients so that Krylov methods reach the best performances in terms of robustness and efficiency.

In a high-fidelity aerostructural optimization context, Zhang and Zingg [1] implemented a robust monolithic solution method for both aerostructural analysis and coupled adjoint problem. A three-dimensional formulation was adopted involving the mesh, the flow and the structural states. The performance of the monolithic method as well as the partitioned one was investigated through a comparative study by varying the level of fluid-structure coupling. For the coupled adjoint problem solution, a GCROT [2] Krylov solver has been used in conjunction with a block Gauss-Seidel preconditioner. The monolithic adjoint solution has been 60 % more efficient than the partitioned one for strong coupling. For weak fluid-structure coupling, the monolithic solution still outperformed the partitioned one with a better efficiency of 40 %. In terms of computational time, the monolithic method showed 50 % to over 60 % faster than the partitioned method. A similar comparative study of both monolithic and partitioned approaches was performed by Kenway et al. [3] except that a Block Jacobi preconditioner was applied to the coupled adjoint system. The aerodynamic and structural block preconditioners were solved by a preconditioned Krylov method (restarted Generalized Minimal RESidual - GMRES [4]) and a direct factorization method respectively. A Flexible Krylov method (e.g FGMRES [5]) has been used for the coupled adjoint system solution. The Common Research Model (CRM) wing-body-tail configuration was sized by considering two critical load cases: $1g$ cruise condition with moderate elastic deformation and a $2.5g$ pull-up with significantly more deflection. For the same memory footprint, the best monolithic solution seems to outperform the best partitioned one by reducing the time to 19% for the $1g$ load and by 29% for the $2.5g$ case. These numerical experiments demonstrate the great benefit of using monolithic approach in the strong coupling case but at the price of a robust preconditioner for the Krylov solver. We note however that both studies only considered inviscid flow modeling. These conclusions about the monolithic solver efficiency might be mitigated by the added stiffness of adjoint system matrices produced by a RANS fluid model associated to a linearized turbulence model.

Although the satisfactory performance of monolithic solvers, advanced strategies that could accelerate the partitioned algorithm using black-box solvers have received less attention. We recall that a partitioned algorithm consists in approximately solving the aerodynamic adjoint block at each fluid-structure iteration resulting in a sequence of adjoint linear systems with varying right-hand sides. As already mentioned, the structural source term that affects the right-hand side of the fluid block has the nice property of rapidly converging to a constant value. The corollary of this property is that after several fluid-structure couplings, the subsequent fluid systems should greatly benefit from recycling spectral information from the previous fluid-structure cycles. At the start of this work, the current partitioned adjoint solver did not take advantage of recycling and at each update of the structural source term the Krylov solver did a cold start from the previous solution.

The principle of deflation is to remove the influence of a system's subspace on the iterative process. This is usually beneficial when directions of certain subspaces hamper convergence. Deflation of an eigenspace can be performed in two ways: the linear system (matrix and right-hand side) is left-multiplied by a projector P, i.e. deflation by projection [6–8], or some eigenvectors are added to the Krylov subspace, i.e. deflation by augmentation [9]. A survey of deflation and augmentations techniques can be found in [10]. A specific type of deflation preconditioning aims at solving a rank-deficient projected system, using a Krylov solver, in a certain subspace outside of the problematic subspace. The solution is then complemented with the solution in the latter subspace. In [8] the author makes the post-correction superfluous by using a projection as right-preconditioner instead. For deflation by augmentation, adding eigenvectors to the Krylov subspace can effectively deflate corresponding eigenvalues from the spectrum because when these directions are included in the solution approximation, the convergence of the Krylov solver continues according to the modified spectrum. The deflation by augmentation led to the well-known FGMRES-DR solver [11] and its extension to inner-outer Krylov solver [12, 13].

Unfortunately, the deflated restarted GMRES framework based on subspace augmentation is not adapted for solving sequences of linear systems [14]. Fortunately, some authors have proposed new strategies in order to reuse information accumulated in previous fluid-structure cycles and use it to accelerate the solution of the next linear system. Krylov subspaces recycling methods seem to be the suitable choice. Historically, De Sturler suggested the Generalized Conjugate Residuals with inner Orthogonalization (GCRO) method [15], an improved version of the recursive GMRES (GMRESR) solver [16] by maintaining an orthogonality condition between the outer and the inner spaces generated by GMRESR. This way, it provides the optimal correction to the solution in a global search space. Later, Parks et al. formulated the GCRO-DR algorithm [14] that combines GCRO and deflation techniques by augmentation introduced by Morgan

[9]. They demonstrated better performances of GCRO-DR compared to the GMRES-DR in a long sequence of linear systems from a fracture mechanics problem. Carvhalo et al. extended GCRO-DR to the flexible case (FGCRO-DR) [17] and they conducted in-depth analysis of both flexible methods. In particular, they showed that both methods can be algebraically equivalent if a certain colinearity condition is satisfied at each cycle. In 2013, Niu et al. introduced Loose GCRO-DR (LGCRO-DR) [18] for improving the convergence of GCRO-DR by recycling both spectral information and approximate error information. The error is defined as the distance between the current iterate and the exact solution of the system. It is not known by definition but a fair approximation to it can be computed. This idea was initially proposed by Baker et al. [19] and mimics the idea behind GMRESR of including approximations to the error in the current approximation space. This error information is interesting since it represents in some sense the previous Krylov space generated in the previous cycle and subsequently discarded. In addition to that, LGCRO-DR is straightforward and economic to implement.

In this paper, we investigate advanced Krylov subspace methods using subspace recycling strategies for accelerating the partitioned solver. More specifically, we compare the performances of flexible GCRO-DR with and without subspace recycling and flexible GMRES-DR for varying degrees of fluid-structure coupling. This work will benefit from the recent achievements to improve efficiency of the fluid adjoint solution by applying nested Krylov subspace methods [13]. The numerical experiments are performed on an aeroelastic configuration of the ONERA M6 fixed wing in transonic viscous flow.

This paper is organized as follows. In section II we briefly recall the theoretical background of aeroelastic and coupled-adjoint equations. The partitioned algorithm is also outlined. Section III is devoted to the description of the GCRO algorithm with adaptations related to variable preconditioning and subspace recycling. The aeroelastic numerical test case is presented in section IV. In section V we detail our preliminary numerical experiments. First, we show that FGMRES-DR and FGCRO-DR are numerically equivalent for solving the fluid adjoint system as demonstrated in [17]. Then, the benefit of adding more physics in the fluid preconditioner is demonstrated for the standard coupled-adjoint partitioned solver. Last, we present preliminary results illustrating the acceleration of the partitioned solver thanks to subspace recycling.

II. Aerostructural adjoint system

A. Aeroelastic equilibrium

Let us denote the state variables of the coupled system $\mathbf{W}$ and $\mathbf{U}$, representing the fluid conservative variables and the structural displacements respectively. At the aeroelastic equilibrium, the state variables and the meshes satisfy the discretized equations of fluid and structural mechanics simultaneously:

$$\begin{cases} \mathbf{R}_a(\mathbf{X}_a, \mathbf{W}, \mathbf{U}) &= \mathbf{0} \\ \mathbf{R}_s(\mathbf{X}_s, \mathbf{W}, \mathbf{U}) &= \mathbf{0} \end{cases} \tag{1}$$

where $\mathbf{R}_a$ is the discrete aerodynamic residual and $\mathbf{R}_s$ the discrete structural residual. These two blocks of equations are coupled through aerodynamic forces $\mathbf{Q}_a$ loading the skin of the structure and the structural displacements $\mathbf{U}$ deforming the fluid mesh. The structural mesh is noted $\mathbf{X}_s$. In the following we introduce two aerodynamic grids $\mathbf{X}_a$ and $\mathbf{X}_{a0}$. $\mathbf{X}_a$ denotes the deformed aerodynamic grid at the aeroelastic equilibrium at the outcome of the aeroelastic analysis. $\mathbf{X}_{a0}$ is called the reference mesh which supports the aerodynamic shape parametrization. Typically for an aircraft design study the reference mesh is chosen as the jig shape or the flight shape in reference nominal flight conditions. The load, displacement and mesh deformation operators then merely depend on $\mathbf{X}_{a0}$ for an aeroelastic or coupled-adjoint analysis. The structural loads $\mathbf{Q}_s$ are obtained with a suitable load transfer technique applied to $\mathbf{Q}_a$ such that

$$\mathbf{Q}_s(\mathbf{Q}_a(\mathbf{W}, \mathbf{X}_a), \mathbf{X}_{a0}, \mathbf{X}_s) = \mathbf{T}^{Q}_{surf}(\mathbf{X}_{a0}, \mathbf{X}_s)\mathbf{Q}_a(\mathbf{W}, \mathbf{X}_a) \tag{2}$$

where $\mathbf{T}^{Q}_{surf}$ represents a linear load transfer operator. The subscript $surf$ stipulates that the associated linear operators or data relate to the fluid-structure interface. The structural displacements alter the fluid grid locations through the relation:

$$\mathbf{X}_a = \mathbf{X}_{a0} + \delta\mathbf{X}_a(\delta\mathbf{X}_{a,surf}, \mathbf{X}_{a0}) = \mathbf{X}_{a0} + \mathbf{T}_{vol}(\mathbf{X}_{a0})\delta\mathbf{X}_{a,surf} \tag{3}$$

with $\mathbf{T}_{vol}(\mathbf{X}_{a0})$ the volume operator performing the deformation of the fluid domain. The vector $\delta\mathbf{X}_{a,surf}$ corresponds to the displacements of the fluid nodes at the fluid-structure interface.

$$\delta\mathbf{X}_{a,surf} = \delta\mathbf{X}_{a,surf}(\mathbf{X}_{a0}, \mathbf{X}_s, \mathbf{U}) = \mathbf{T}^U_{surf}(\mathbf{X}_{a0}, \mathbf{X}_s)\mathbf{U} \tag{4}$$

where $\mathbf{T}^U_{surf}(\mathbf{X}_{a0}, \mathbf{X}_s)$ represents a linear displacement transfer operator.

B. Partitioned strategy for the coupled adjoint system

Let us consider a scalar aeroelastic objective function $J(\mathbf{W}, \mathbf{X}_a, \mathbf{X}_s)$ and a design parameter p. One way to obtain the coupled adjoint equations is to formulate an augmented objective function by adding the total variation of the residuals $\mathbf{R}_s$ and $\mathbf{R}_a$ to the total derivative dJ/dp. More specifically, we define $d\tilde{J}/dp$ as

$$\frac{d\tilde{J}}{dp} = \frac{dJ}{dp} + \lambda_a^T \frac{d\mathbf{R}_a}{dp} + \lambda_s^T \frac{d\mathbf{R}_s}{dp} \tag{5}$$

where

$$\frac{dJ}{dp} = \frac{\partial J}{\partial \mathbf{W}}\frac{d\mathbf{W}}{dp} + \frac{\partial J}{\partial \mathbf{X}_a}\frac{d\mathbf{X}_a}{dp} + \frac{\partial J}{\partial \mathbf{X}_s}\frac{d\mathbf{X}_s}{dp} + \frac{\partial J}{\partial \mathbf{U}}\frac{\partial \mathbf{U}}{\partial p}, \tag{6}$$

In Eq. (5) the total variations of residuals are exactly zero since they represent constraints related to the satisfaction of the equilibrium equations at the outcome of the aeroelastic analysis. For simplicity, we restrict here to the specific case of a shape design parameter not affecting the structural geometry nor the structural stiffness. In addition, the explicit dependency of the objective function with respect to the structural states is dropped, i.e., we consider only derivatives of aerodynamic coefficients. The full derivation for the general case can be found in [20, 21]. It is worth to mention that these assumptions do not lead to any loss of generality of the work presented in this paper since we focus on solution techniques for the adjoint system. As $\mathbf{X}_s$ does not depend on the design parameter p, we have $d\mathbf{X}_s/dp = \mathbf{0}$. Under the same assumption we also have $d\mathbf{K}/dp = \mathbf{0}$. After some algebra manipulation we end up with the following expression for $d\tilde{J}/dp$ in which the total derivatives $d\mathbf{W}/dp$ and $d\mathbf{U}/dp$ have been factored out:

$$\begin{aligned}\frac{d\tilde{J}}{dp} = &\left(\frac{\partial J}{\partial \mathbf{W}} + \lambda_a^T \frac{\partial \mathbf{R}_a}{\partial \mathbf{W}} - \lambda_s^T \mathbf{C}\right)\frac{d\mathbf{W}}{dp} + \left(\frac{\partial J}{\partial \mathbf{X}_a}\mathbf{A} + \lambda_a^T \frac{\partial \mathbf{R}_a}{\partial \mathbf{X}_a}\mathbf{A} + \lambda_s^T(\mathbf{K} - \mathbf{D})\right)\frac{d\mathbf{U}}{dp} \\ &+ \left(\frac{\partial J}{\partial \mathbf{X}_a} + \lambda_a^T \frac{\partial \mathbf{R}_a}{\partial \mathbf{X}_a}\right)\mathbf{B}\frac{d\mathbf{X}_{a0}}{dp} - \lambda_s^T \mathbf{E}\frac{d\mathbf{X}_{a0}}{dp}\end{aligned} \tag{7}$$

Constant matrices $\mathbf{A}$ to $\mathbf{E}$ are defined analytically with the following formulas (see [20]):

$$\mathbf{A} = \mathbf{T}_{vol}\mathbf{T}^U_{surf} \tag{8}$$

$$\mathbf{B} = \frac{\partial \mathbf{X}_a}{\partial \mathbf{X}_{a0}} = \mathbf{I} + \frac{\partial \mathbf{A}}{\partial \mathbf{X}_{a0}}\mathbf{U} \tag{9}$$

$$\mathbf{C} = \mathbf{T}^Q_{surf}\frac{\partial \mathbf{Q}_a}{\partial \mathbf{W}} \tag{10}$$

$$\mathbf{D} = \mathbf{T}^Q_{surf}\frac{\partial \mathbf{Q}_a}{\partial \mathbf{X}_a}\mathbf{T}^U_{surf} \tag{11}$$

$$\mathbf{E} = \mathbf{T}^Q_{surf}\frac{\partial \mathbf{Q}_a}{\partial \mathbf{X}_a}\mathbf{B} + \frac{\partial \mathbf{Q}_s}{\partial \mathbf{X}_{a0}} \tag{12}$$

The coupled adjoint linear system is obtained by canceling factors related to $d\mathbf{W}/dp$ and $d\mathbf{U}/dp$ in Eq. (7) to give

$$\begin{bmatrix} \left[\frac{\partial \mathbf{R}_a}{\partial \mathbf{W}}\right]^T & -\mathbf{C}^T \\ \mathbf{A}^T\left[\frac{\partial \mathbf{R}_a}{\partial \mathbf{X}_a}\right]^T & \mathbf{K}^T - \mathbf{D}^T \end{bmatrix} \begin{bmatrix} \lambda_a \\ \lambda_s \end{bmatrix} = \begin{bmatrix} -\left[\frac{\partial J}{\partial \mathbf{W}}\right]^T \\ -\mathbf{A}^T\left[\frac{\partial J}{\partial \mathbf{X}_a}\right]^T \end{bmatrix} \tag{13}$$

The process for solving the adjoint system follows an iterative block scheme. Algorithm 1 details the Linear Block Gauss Seidel (LBGS) scheme applied for the solution of system (13). In this derivation, we use the structural flexibility $\mathbf{S}$ which is a small reduced matrix relating the set of structural forces to the set of structural displacments pertaining to the fluid-structure coupling. The relaxation factor θ_s has been introduced on the adjoint vector λ_s. Assuming that the coupled system is solved to machine accuracy, the total derivative reconstruction is given by

$$\frac{dJ}{dp} = \left(\frac{\partial J}{\partial \mathbf{X}_a} + \lambda_a^T \frac{\partial \mathbf{R}_a}{\partial \mathbf{X}_a} \right) \mathbf{B} \frac{d\mathbf{X}_{a0}}{dp} - \lambda_s^T \mathbf{E} \frac{d\mathbf{X}_{a0}}{dp} \tag{14}$$

In the expression above the computation of the product of the geometrical sensitivities with the matrix $\mathbf{B}$ is not trivial. If one already has at hand a linearized version of the operator $\mathbf{A}$, i.e. of $\mathbf{T}^U_{surf}$ and $\mathbf{T}_{vol}$, it can be applied to $d\mathbf{X}_{a0}/dp$ as many times as the number of design variables. This is the most straightforward manner but the benefit of the adjoint formulation is then mitigated by the cost of the gradient assembly. The other way is to transpose the first term in the right-hand side of Eq. (14) and compute products like $[\partial \mathbf{T}^U_{surf}/\partial \mathbf{X}_{a0}]^T \mathbf{v}$ and $[\partial \mathbf{T}_{vol}/\partial \mathbf{X}_{a0}]^T \mathbf{v}$, where $\mathbf{v}$ has fluid grid size. We call this mode the geometrical adjoint of the mesh deformation and displacement transfer operators. These two modes of gradient assembly have been implemented in the coupled-adjoint module of the elsA software.

Algorithm 1 Partitioned LBGS strategy for coupled-adjoint solution

Initialize: $\mathbf{U}, \mathbf{W}\ \mathbf{X}_a, \mathbf{X}_s, \mathbf{X}_{a0}, \lambda_a^0, \lambda_s^0, \theta_s, tol$

if $\lambda_s^0 \neq \mathbf{0}$ **then**

  $\mathbf{RHS}_{stru} \leftarrow \left(\mathbf{T}^Q_{surf} \frac{\partial \mathbf{Q}_a}{\partial \mathbf{W}} \right)^T \mathbf{S}^T \lambda_s^0$

else

  $\mathbf{RHS}_{stru} \leftarrow \mathbf{0}$

end if

if $\lambda_a^0 = \mathbf{0}$ **then**

  $\left[\frac{\partial \mathbf{R}_a}{\partial \mathbf{W}} \right]^T \lambda_a^0 = - \left[\frac{\partial J}{\partial \mathbf{W}} \right]^T + \mathbf{RHS}_{stru}$ ▷ Approximate solution of the fluid adjoint problem

end if

for $k \leftarrow 1, n_{cpl}$ **do**

  $A_{Xs,surf} \leftarrow \left(\frac{\partial \mathbf{Q}_a}{\partial \mathbf{X}_{a,surf}} \right)^T \left(\mathbf{T}^Q_{surf} \right)^T \mathbf{S}^T \lambda_s^{k-1}$ ▷ Structural geometric adjoint

  $A_{Xa} \leftarrow - \left([\lambda_a^{k-1}]^T \frac{\partial \mathbf{R}_a}{\partial \mathbf{X}_a} + \frac{\partial J}{\partial \mathbf{X}_a} \right)^T$ ▷ Aerodynamic geometric adjoint

  $A_{Xa,surf} \leftarrow (\mathbf{T}_{vol})^T \mathbf{A}_{Xa}$ ▷ Mesh deformation adjoint

  $\lambda_s^k \leftarrow (\mathbf{T}^U_{surf})^T (A_{Xs,surf} + A_{Xa,surf})$ ▷ Structural adjoint vector

  $\lambda_s^k \leftarrow (1 - \theta_s) \lambda_s^{k-1} + \theta_s \lambda_s^k$ ▷ Relaxation

  $\mathbf{RHS}_{stru} \leftarrow \left(\mathbf{T}^Q_{surf} \frac{\partial \mathbf{Q}_a}{\partial \mathbf{W}} \right)^T \mathbf{S}^T \lambda_s^k$ ▷ Update structural rhs

  $\left[\frac{\partial \mathbf{R}_a}{\partial \mathbf{W}} \right]^T \lambda_a^k = - \left[\frac{\partial J}{\partial \mathbf{W}} \right]^T + \mathbf{RHS}_{stru}$ ▷ Approximate solution of the fluid adjoint problem

end for

$\frac{dJ}{dp} \leftarrow \left(\frac{\partial J}{\partial \mathbf{X}_a} + \lambda_a^T \frac{\partial \mathbf{R}_a}{\partial \mathbf{X}_a} \right) \mathbf{B} \frac{d\mathbf{X}_{a0}}{dp} - \lambda_s^T \mathbf{E} \frac{d\mathbf{X}_{a0}}{dp}$ ▷ Objective function gradient assembly

III. Recycling strategies for sequences of linear systems

As already explained, the current implementation of our partitioned solver does not take advantage of spectral information produced during the Krylov solver instances applied to the sequence of preceding right-hand sides. An advanced inner-outer GMRES-DR solver is used for the approximate solution of the fluid block between consecutive fluid-structure couplings but a simple restart is performed when the structural source term is updated. Unfortunately, deflated restarting based on subspace augmentation by appending approximate Ritz vectors to the Krylov subspace is not suitable for sequences of linear systems. Indeed, if the Ritz vectors are obtained from a previous linear system with another matrix or even another right-hand side, the concatenation of the recycled subspace to the current Krylov subspace does not form a Krylov subspace for the current problem. It is necessary to introduce a new Krylov solver that uses recycling of any given subspace without restriction. A famous one is the generalized conjugate residual with inner orthogonalization (GCRO). It belongs to the family of inner-outer methods [22] where the outer method is based on GCR, a minimum residual norm method proposed by Eisenstat, Elman and Schultz [23]. The inner solver is GMRES applied to a projected system matrix.

A. Flexible Generalized Conjugate Residual with inner Orthogonalization and Deflated Restarting: FGCRO-DR

As previously mentioned, we focus on the solution of a sequence of linear systems with a varying right-hand side only. This section briefly introduces the Flexible GCRO algorithm with deflated restarting. In this framework, deflation can reuse spectral information from a previous cycle or from a previous linear system. We start by recalling the original formulation of the generalized conjugate residual method (GCR)[23]. The idea is to introduce the concept of optimality for the solution residual. We want to solve

$$Ax^{(i)} = b^{(i)}, \quad i = 1, 2, \ldots \tag{15}$$

where $A \in \mathbb{R}^{n\times n}$ and $b^{(i)} \in \mathbb{R}^n$ changes from one system to the next.

The GCR method relies on a given full-rank matrix $Z_k \in \mathbb{R}^{n\times k}$ and an orthonormal matrix $C_k \in \mathbb{R}^{n\times k}$ as the image of Z_k by A satisfying the relations

$$AZ_k = C_k, \tag{16}$$

$$C_k^T C_k = I_k. \tag{17}$$

For the sake of understanding, suppose that we want to solve the first system of the sequence of linear systems, i.e., $i = 1$. Given an initial guess x_0^1, the principle is to compute an approximation to the solution $x^1 \in x_0^1 + \text{range}(Z_k)$ that minimizes the corresponding residual norm in the approximation space range(Z_k). More precisely, GCR solves the following minimization problem

$$x^{(1)} = \underset{x \in x_0^{(1)} + \text{range}(Z_k)}{\text{argmin}} \|b^{(1)} - Ax\|_2, \tag{18}$$

The optimal solution of (18) over the subspace $x_0^{(1)} + \text{range}(Z_k)$ is defined by

$$x^{(1)} = x_0^{(1)} + Z_k C_k^T r_0^{(1)} \tag{19}$$

Consequently, the corresponding residual vector satisfies

$$r_k^{(1)} = b^{(1)} - Ax^{(1)} = (I - AZ_k C_k^T) r_0^{(1)} = (I - C_k C_k^T) r_0^{(1)}, \quad r_k^{(1)} \perp \text{range}(C_k). \tag{20}$$

The orthogonality of the residual $r_k^{(1)}$ to the subspace $A\mathcal{K}_k(A, r_0^{(1)})$ spanned by the columns of C_k is known as the optimality property of the residual. In practice, GCR is not considered as a means to solve the linear system. It is replaced by a GMRES solver which rather computes an implicit representation of the matrices Z_k and C_k [2].

Given an orthonormal basis C_k of an outer subspace and the corresponding residual $r_k = (I - C_k C_k^T) r_0$ after k steps of GCR, GCRO(m) obtains the next iterates r_{k+1} and x_{k+1} by performing m steps of GMRES applied to the projected operator $A_{C_k} = (I - C_k C_k^T)A$, thereby maintaining optimality of the inner residual to the outer space. Let $V_{m+1} \in \mathbb{R}^{n\times(m+1)}$ be an orthonormal basis for $\mathcal{K}_{m+1}(A_{C_k}, r_k)$ with $v_1 = r_k/\|r_k\|$, if we consider a flexible inner GMRES we have the following Arnoldi relation:

$$A_{C_k} Z_m = (I - C_k C_k^T) A Z_m = V_{m+1} \bar{H}_m \quad \text{with} \quad \bar{H}_m = V_{m+1}^T A_{C_k} Z_m \tag{21}$$

This equation can be expanded as

$$AZ_m = C_k B_m + V_{m+1} \bar{H}_m \quad \text{with} \quad B_m = C_k^T A Z_m \tag{22}$$

At the end of the inner GMRES cycle, the residual is expressed as

$$r_{k+1} = r_k - A_{C_k} Z_m y = r_k - V_{m+1} \bar{H}_m y_m \tag{23}$$

with, using $\bar{H}_m = \bar{Q}_m R_m$, y_m such that

$$y_m = \underset{y \in \mathbb{R}^m}{\operatorname{argmin}} \|r_k - A_{C_k} Z_m y\|_2 = R_m^{-1} \bar{Q}_m^T \|r_k\|_2 e_1 \tag{24}$$

In the above Arnoldi process (22) of the inner GMRES, the vectors Az_i are first orthogonalized against C_k, thus constructing V_{m+1} such that $C_k^T V_{m+1} = 0$. So far, we have not yet explained how the matrix C_k is built in the context of subspace recycling. First, we slightly change the notation in (21) and (22) to point out that we perform $m - k$ steps (instead of m) of inner FGMRES. This produces the Arnoldi relation

$$A_{C_k} Z_m = (I - C_k C_k^T) A Z_{m-k} = V_{m-k+1} \bar{H}_{m-k} \tag{25}$$

Combining (16) and (25), where C_k is renamed V_k for better readability, we have

$$A[Z_k \quad Z_{m-k}] = [V_k \quad V_{m-k+1}] \begin{bmatrix} I_k & B_{m-k} \\ 0 & \bar{H}_{m-k} \end{bmatrix} \quad \text{with} \quad B_{m-k} = V_k^T A Z_{m-k} \tag{26}$$

From now on, we define $Z_m = [Z_k \quad Z_{m-k}]$ and $V_{m+1} = [V_k \quad V_{m-k+1}]$ leading to the following compact expression of (26)

$$AZ_m = V_{m+1} \bar{H}_m. \tag{27}$$

We have now completed the process for the first linear system which formulates the solution from the complementary subspaces V_k and V_{m-k+1}. At this stage, we introduce deflated restarting within the FGCRO framework to build the outer space V_k for the next linear system. The deflation subspace is spanned by the k harmonic Ritz vectors corresponding to harmonic Ritz values of smallest magnitude. This leads to the FGCRO-DR algorithm that is a combination of FGCRO and FGMRES-DR (see Algorithm 2). To compute the harmonic Ritz vectors, we solve the $m \times m$ generalized eigenvalue problem defined in step 28 of Algorithm 2. We then form the matrix $P_k \in \mathbb{R}^{m \times k}$ which stores the retained eigenvectors. This matrix is expanded by Z_m to form the basis of the deflation subspace $Z_m P_k$. Then, the image of the deflation space by A gives

$$AZ_m P_k = V_{m+1} \bar{H}_m P_k \tag{28}$$

then orthonormalized,

$$V_k = V_{m+1} \bar{Q}_m \tag{29}$$

where V_{m+1} has been computed at the end of the $(i-1)$st cycle and $\bar{Q}_m$ stands for the unitary factor of the QR factorization of $\bar{H}_m P_k = \bar{Q}_m R_k$. The relation that holds at this stage is:

$$AZ_m P_k R_k^{-1} = V_k \tag{30}$$

From (16), we deduce that $Z_k = Z_m P_k R_k^{-1}$. Knowing V_k and Z_k, the solution and the residual can be updated with

$$x_1^{(i+1)} = x_0^{(i+1)} + Z_k V_k^T r_0^{(i+1)} \tag{31}$$

$$r_1^{(i+1)} = (I - V_k V_k^T) r_0^{(i+1)} \tag{32}$$

and we can initiate the Arnoldi process in (26) for the next system.

IV. ONERA-M6 wing aeroelastic analysis

In this study the numerical experiments have been performed with the well known ONERA-M6 fixed wing configuration which has been extensively used for CFD solvers validation in transonic flow conditions. In this work we use the RANS solver provided by elsA for the steady rigid and aeroelastic analyses [24, 25]. The elsA adjoint and coupled-adjoint solvers have also been used for the computation of rigid and flexible derivatives. The latest improvements to the Krylov solvers for the solution of the adjoint linear system are described in [13]. A multi-block structured mesh featuring a C-H topology is used (Fig. 1). It consists of 3.8 million grids divided into 42 blocks. The flight conditions are a free-stream Mach number of 0.84 at an angle of attack of 3.06 degrees. The convective fluxes are discretized by an upwind Roe scheme associated to a Monotonic Upstream-centered Scheme for Conservation Laws (MUSCL) reconstruction and a Van Albada flux limiter. The one-equation Spalart-Allmaras turbulence model is selected. The surface contours in the bottom plot of Fig. 2 below show typical results for the ONERA-M6 wing. The pressure coefficient contours identify a lambda-shock along the mid-chord of the wing. For the aeroelastic analysis a simple but realistic finite element model has been designed (Fig. 1). The stiffness of this model can be easily tuned to get stronger or weaker fluid-structure interaction. The pressure coefficient contours at the aeroelastic equilibrium are plotted in the upper part of Fig. 2 and can be compared to the rigid contours. The maximum vertical displacement is 0.14 meters corresponding to 11.7 % of the wing span. To get a better insight of the effect of flexibility on the pressure distribution, we report in Fig. 3 the Cp distributions for two sections at $y = 0.60$ m and $y = 1.12$ m. The vertical displacement distributions associated to the front and rear spars as well as the twist increment distribution are plotted in Fig. 4. The rigid analysis results in a lift coefficient $C_L = 0.27$ whereas the aeroelastic analysis, at the same angle of attack, results in a lower lift coefficient $C_L = 0.23$.

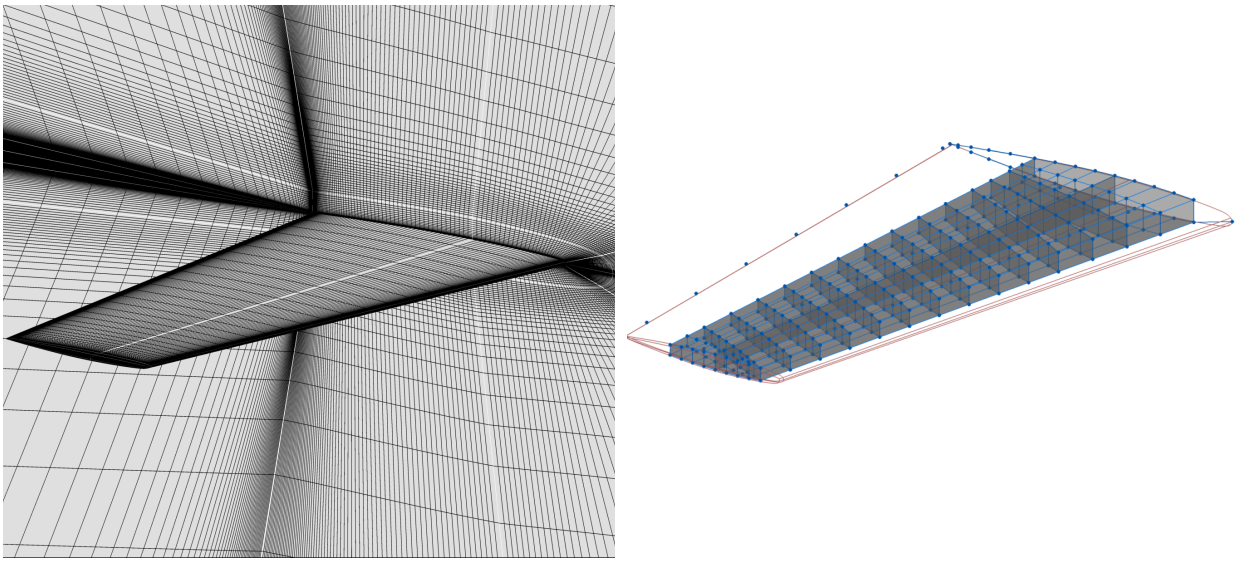

Fig. 1 M6 wing aeroelastic model: 42 block-structured RANS CFD mesh and FEM internal layout.

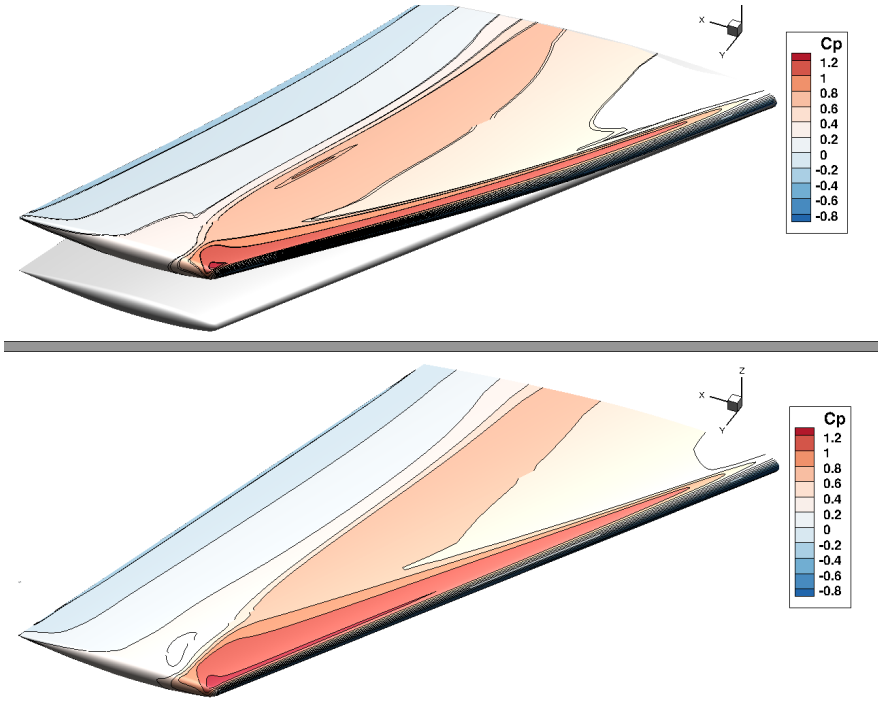

Fig. 2 Pressure coefficient contour plots for the rigid and aeroelastic steady flows.

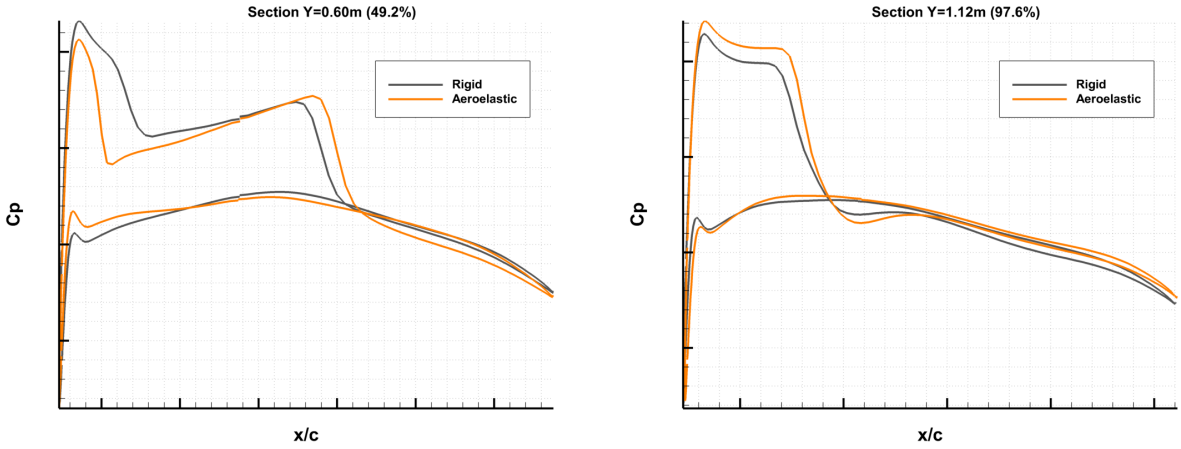

Fig. 3 Comparison of rigid and aeroelastic pressure coefficient section plots at y=0.60m and y=1.12m.

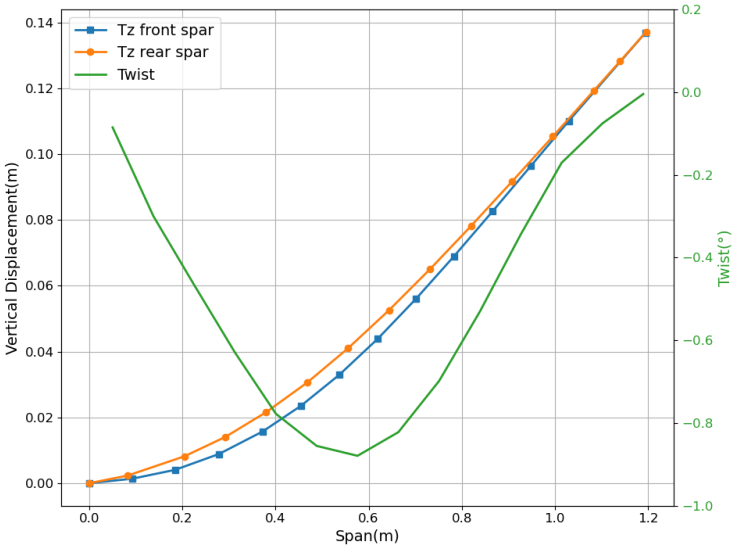

Fig. 4 Vertical displacement and twist increment distribution at aeroelastic equilibrium.

V. Numerical experiments

A. Preconditioning

We are interested in the solution of linear systems of adjoint equations for a finite volume discretization on multi-block structured meshes. As already mentioned, all computations are performed using a RANS flow model and a fully linearized one-equation Spalart-Allmaras turbulence model. The spatial discretization relies on a second-order upwind Roe scheme. The associated system of equations is usually very stiff as it exhibits terms of very different amplitudes, depending on the location in the fluid domain, e.g., in the flow wake, near the wall surface or at the far boundary. Some recent work by the authors [13] demonstrated the numerical robustness of the inner-outer GMRES solvers for such difficult problems. This class of solvers embeds two level of preconditioning. The outer preconditioner is a standard GMRES solver pre-conditioned itself by an inner stationary operator. The GMRES preconditioner offers the desirable property of addressing the global eigenspectrum of the entire domain. This also makes the inner-outer solver less sensitive to the quality of the second-level preconditioner.

Two stationary preconditioning strategies are considered for the inner GMRES. The first one consists in a block version of a standard block Lower-Upper Symmetric Gauss-Seidel (LU-SGS) iterative solver [26]. LU-SGS is a particular case of LU-SSOR with a unity relaxation factor. The reader should refer to [27] for further details about LU-SGS as an iterative solver for the implicit phase of Backward-Euler Newton steady solvers. In the current implementation, LU-SGS is combined with a domain decomposition method. More specifically, the globalization of the preconditioner is achieved with a Restricted Additive Schwarz method [28].

LU-SGS is applied to a first-order diagonally dominant upwind approximation of the flux Jacobian matrix inspired by [29]. This operator is based on a first-order spatial discretization of the convective and of the viscous fluxes using a simplifying thin layer assumption [27]. This strategy leads to a very compact stencil for the preconditioning matrix which will be denoted as $\mathbf{J}_{O1}^{APP}$ in the remainder of this paper. The second one is a Block Incomplete Lower Upper (BILU(k)) factorization applied to either an approximate or exact flux Jacobian matrix. For the so-called first-order exact Jacobian matrix $\mathbf{J}_{O1}^{EX}$, a first-order spatial Roe scheme is used for the discretization of the mean-flow convective fluxes and a 5-point corrected centered discretization scheme is used for the diffusive fluxes. The BILU(k) preconditioner will be applied either to the first-order approximate Jacobian matrix $\mathbf{J}_{O1}^{APP}$, or to the first-order exact Jacobian matrix $\mathbf{J}_{O1}^{EX}$. About memory footprint, $\mathbf{J}_{O1}^{EX}$ has a 9-point stencil in 2D whereas a 5-point stencil is associated with $\mathbf{J}_{O1}^{APP}$. In 3D, we have a 7-point stencil for $\mathbf{J}_{O1}^{APP}$ and a stencil of 19 points for $\mathbf{J}_{O1}^{EX}$. Consequently, a better robustness is achieved but at the price of about twice the storage for $\mathbf{J}_{O1}^{EX}$ compared to $\mathbf{J}_{O1}^{APP}$. For a better understanding, we have reproduced the 3D stencils in Fig. 5.

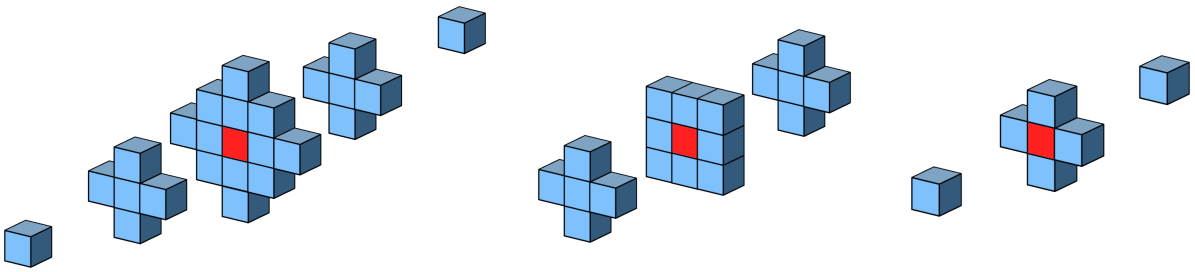

a) Exact 2^{nd} order stencil (size 25) b) Exact 1^{st} order stencil (size 19) c) Approx. 1^{st} order stencil (size 7)

Fig. 5 Stencils of the various Jacobian matrices.

In order to efficiently tackle the different scales in the system Jacobian matrix, an additional level of preconditioning is considered in the present work. We perform a block-diagonal left-right scaling, i.e., row-column scaling, of the system matrix leading to a modified system of the form

$$D_L^{-1} A D_U^{-1} M^{-1} t = D_L^{-1} b, \tag{33}$$

where D_L^{-1} and D_U^{-1} are lower and upper inverses of the LU factors of the block-diagonal of matrix A. This leads to a scaled system matrix with unit diagonal. The preconditioner is then based on an approximate BILU factorization of $D_L^{-1} A D_U^{-1}$. Even if the scaling is performed externally, i.e., the approximate Jacobian matrix is left/right pre-multiplied before applying BILU, the condition number of the BILU factors is greatly reduced and the rounding errors during the factorization process are smaller. About products involving the system matrix A, the operation is performed on-the-fly as the matrix is not stored.

B. Performance of recycling Krylov solvers for aerodynamic adjoint systems

Having coded our new FGCRO-DR solver, the primal objective in this section is to assess its performances and compare to FGMRES-FR for the solution of a fluid adjoint problem. This exercise will also confirm the expected algebraic equivalence that exists between FGCRO-DR and FGMRES-DR and proved by Carvalho et al. in [17]. Since we solve a single linear adjoint system, the subspace recycling step (step 2 to step 4 of Algorithm 2 in Appendix A) is skipped. Actually, Fig.6 illustrates an essential numerical aspect pertaining to the FGCRO-DR method. More precisely, FGCRO-DR seems to be more sensitive than FGMRES-DR to the rounding errors during the orthogonalization process. Indeed, for FGCRO-DR we observe a stagnation of the relative residual after 10^{-8} when a single step of orthogonalization is applied, even for a moderate size Krylov subspace (FGCRO-DR(30,10,10) with nMGS=1).

The best way to restore convergence in a reasonable number of iterations is to perform two steps of orthogonalization (FGCRO-DR(30,10,10) with nMGS=2). The main reason that FGCRO-DR suffers from rounding errors during the orthogonalization process comes essentially from the two-step Arnoldi process. Indeed, maintaining the orthogonalization condition between the recycled space V_k and the inner Krylov space V_{m-k} during the Arnoldi process is challenging for our stiff numerical problem. To point out the two-step orthogonalization process, (25) can be rewritten as

$$A Z_{m-k} = V_k B_{m-k} + V_{m-k} \bar{H}_{m-k} \tag{34}$$

In Fig.6 we see that the theoretical numerical equivalence, in exact arithmetics, between FGCRO-DR and FGMRES-DR is demonstrated.

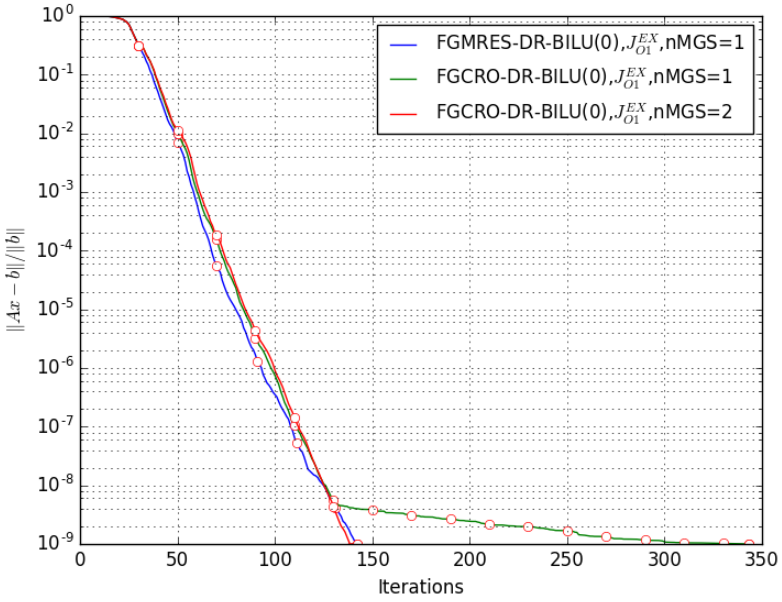

Fig. 6 Comparison of the numerical performances between FGCRO-DR and FGMRES-DR. Impact of the orthogonalization step on the FGCRO-DR solver: one step orthogonalization (nMGS=1) deteriorates the convergence after 10^{-8} while a re-orthogonalization restores the convergence (nMGS=2). We recall the numerical parameters as follows: $m = 30$, $m_i = 10$, $k = 10$.

As we are dealing with RANS flow equations associated to a fully linearized turbulence model, the Jacobian operators may be poorly conditioned in certain areas of the fluid domain. To mitigate this issue, we have considered a block diagonal scaling of the system matrix as defined in (33). Fig.7 illustrates the benefit of this scaling. Indeed, scaling has a strong impact on the stagnation observed at the beginning of the convergence by shrinking the plateau.

Also, true and least-squares relative residuals match more accurately even at tiny convergence levels, i.e., 10^{-9} in this case. Scaling seems then to be a good alternative to re-orthogonalization, and also a cheaper one. In Fig.7 we notice a slightly lower convergence rate for the scaled problem so that both scaled and original problems converge in approximately the same number of iterations. However, in the context of a partitioned solution strategy, we seek for an early decrease of the fluid residual and, in this sense, scaling is very attractive.

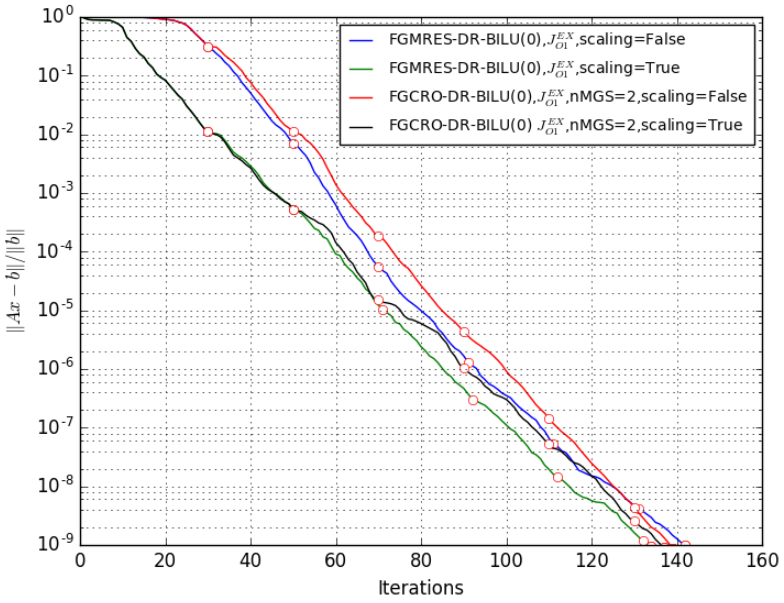

Fig. 7 Adjoint relative residual norm convergence history of FGMRES-DR and FGCRO-DR. Impact of the scaling on the convergence. We recall the numerical parameters of solvers as follows: $m = 30$**,** $m_i = 10$**,** $k = 10$**.**

C. Performance of the current partitioned solver for coupled-adjoint systems

This subsection illustrates the impact of preconditioning strategies applied to the fluid adjoint problem that improve the current partitioned solver in terms of robustness and computational time. We recall that an inner-outer Krylov solver is used in this work and then the selection of the preconditioner of the inner Krylov solver is of interest. The preconditioning strategies have been presented in section V.A.

Considering the relative size of the stencils associated to the exact first-order and the approximate first-order Jacobian matrices, one needs roughly three times more storage for $\mathbf{J}_{O1}^{EX}$ compared to $\mathbf{J}_{O1}^{APP}$ but resulting in a better efficiency as we can see in Fig. 8. Indeed, BILU(0) applied to $\mathbf{J}_{O1}^{EX}$ outperforms LU-SGS with a gain of 45% in terms of iterations. Moreover, the scaling has the ability to enhance the convergence of the coupled-adjoint system by shrinking the plateau that appears at the beginning of each fluid-structure coupling (see Fig.9). From these numerical experiments, we conclude that the best option for performance of the partitioned solver is to associate a diagonal scaling to the BILU(0) factorization of the first-order exact Jacobian matrix. We also recall that the flow equations are systematically physically non-dimensionalized in the first place.

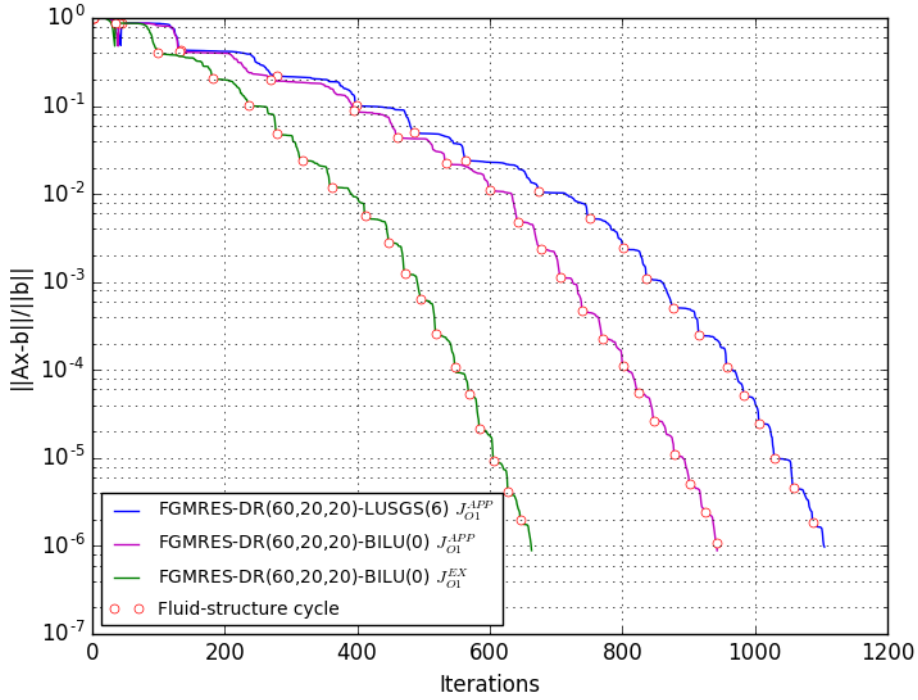

Fig. 8 Adjoint relative residual norm convergence history of FGMRES-DR. Impact of various preconditioners.

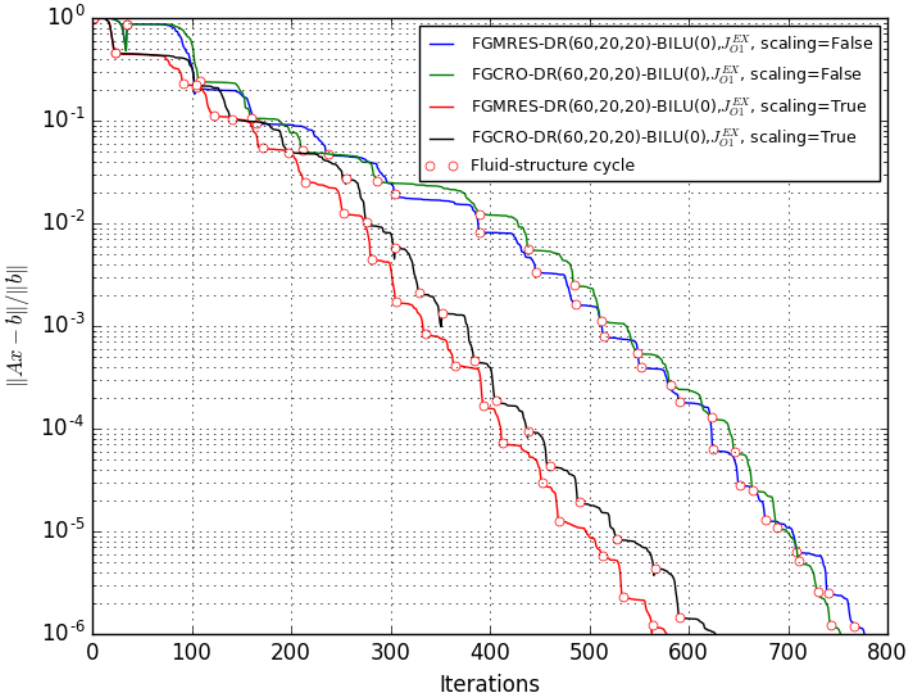

Fig. 9 Adjoint relative residual norm convergence history for FGMRES-DR and FGCRO-DR. Impact of scaling on the convergence. The plateau that appears after each fluid-structure coupling is dropped thanks to the scaling effect.

D. Performance of recycling strategies for aerostructural adjoint systems

The results presented in this section are associated to an approximate linearization using the constant eddy viscosity assumption. However, we point out that the conclusions can be readily transposed to a fully linearized system of equations.

Our objective is to propose simple criteria to assess whether recycling can be activated. The first criterion consists in monitoring the norm of the right-hand side of the fluid block. Indeed, recycling is expected to be efficient for slowly varying systems. As we are solving a coupled problem, we know that after a given number of fluid-structure (F/S) couplings, the structural source term stabilizes. This stabilization of the structural part means that the aerostructural adjoint system solution reduces to an aerodynamic adjoint system only.

In Fig. 10, we plot the assembled right-hand side (red curve) which sums the structural one (RHS_{stru}, blue curve) and the aerodynamic one (RHS_{aero}, green curve). We recall that RHS_{aero} is constant during the solution process and only RHS_{stru} varies. We observe in Fig. 10 that RHS_{stru} varies significantly during the first fluid-structure cycles and stabilizes after six couplings.

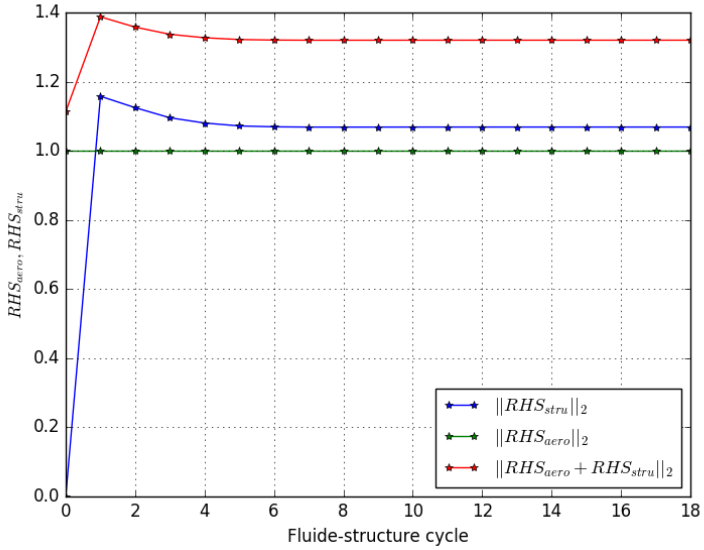

Fig. 10 Variation of structural and aerodynamic right-hand sides with respect to F/S cycles. Residual norms are scaled by the norm of the fluid right-hand side.

Fig. 11 clearly shows the benefit of subspace recycling on the relative residual convergence of the partitioned solver. As just mentioned, the recycling strategy is relevant after six fluid-structure cycles. More precisely, we have considered two cases: the first one triggers the recycling strategy from the fifth cycle and the second one from the sixth cycle. Obviously, recycling is not relevant before six couplings.

Fig. 12 proposes another criterion for the activation of the recycling process. More precisely, at the beginning of the ith fluid-structure cycle, we compare the norms of the initial relative residual $r_0^{(i)}$ and that of the relative residual projected orthogonally to the recycled subspace $V_k = [v_1, ..., v_k]$ denoted as $r_1^{(i)} = (I - V_k V_k^T) r_0^{(i)}$. The ratio $\|r_1^{(i)}\|/\|r_0^{(i)}\|$ can be interpreted as a measure of the degree of colinearity between the two residual vectors. Thus, a value close to unity indicates that the previous approximation V_k of the invariant subspace of the system matrix is relevant for recycling. By the way, as fluid-structure cycles proceed, we expect this approximation to get closer to the true invariant subspace. One interesting way to assess the approximation error is to monitor the distance between subspaces $V_k^{(i)}$ and $V_k^{(i+1)}$. This distance is expressed as $dist(V_k^{(i)}, V_k^{(i+1)}) = \|(I - V_k^{(i)}[V_k^{(i)}]^T)V_k^{(i+1)}[V_k^{(i+1)}]^T\|$.

Fig. 11 reflects these considerations. The convergence of the FGMRES-DR solver is reported for reference. As expected, it is directly comparable to the FGCRO-DR curve without recycling. Turning to the convergence of FGCRO-DR with recycling from the fifth F/S cycle, we observe a degradation of convergence. Clearly, the recycled space is not sufficiently informative to improve convergence during the subsequent cycles. However, recycling at the next cycle dramatically speeds up convergence.

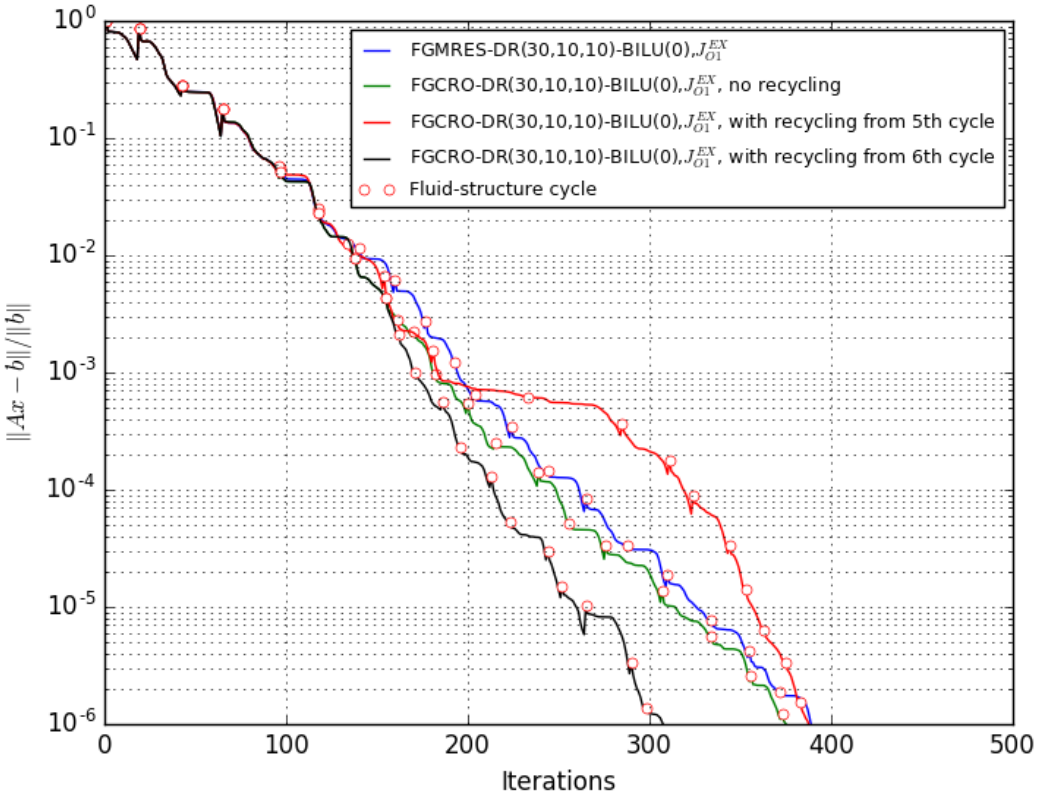

Fig. 11 Relative residual convergence history for both FGCOR-DR and FGMRES-DR. Impact of recycling on the convergence.

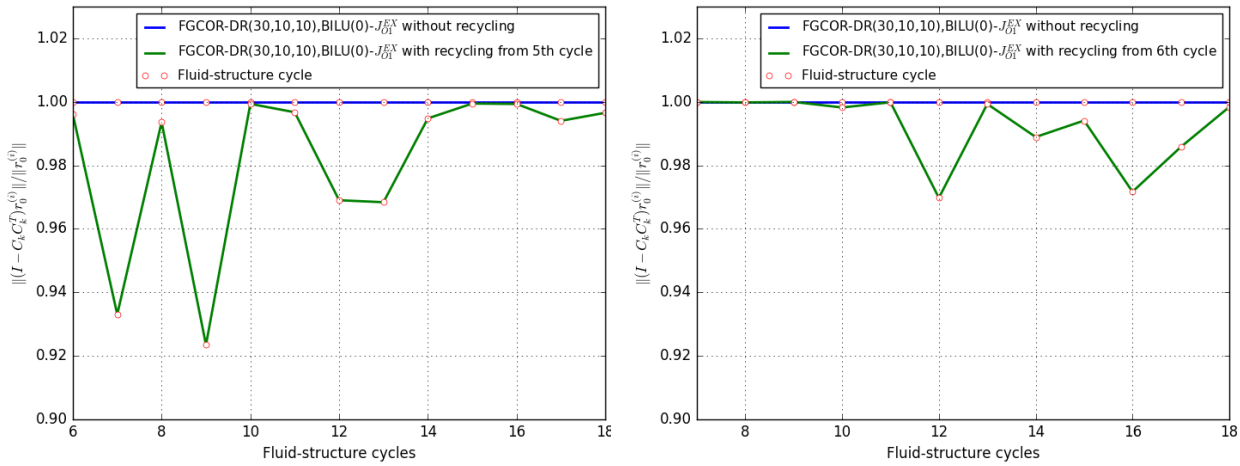

(a) **(b)**

Fig. 12 Contribution of the recycled subspace to the next system solution: (a) corresponds to contribution from the 6th cycle and (b) corresponds from the 7th cycle.

E. Deflation strategies based on error approximation

As part of the GCRO-DR solver, Niu et al. [18] proposed to recycle both the approximate error and the harmonic vectors corresponding to the smallest harmonic Ritz values in magnitude. Let first recall the definition of the error approximation. If we denote x_s and x_{s-1} the approximate solutions obtained at the end of cycle (s) and cycle $(s-1)$ of GMRES(m), and $\hat{x}$ the exact solution, then the approximate error is defined as $z_s = x_s - x_{s-1}$. It represents the correction to x_{s-1} selected from $\mathcal{K}_m(A, r_{s-1})$. The specific choice of z_s for recycling is interesting since it represents the best available approximation of $A^{-1}r_{s-1}$. Indeed, De Sturler demonstrated that the true error vector $e_s = \hat{x} - x_s$ is the optimal solution of the linear system $Ae = r_{s-1}$. Therefore, deflating the error approximation z_s enables the GCRO-DR solver to compute a new search direction thus avoiding the z_s-direction in the subsequent cycles. This strategy leads to a new solver which is denoted as LGCRO-DR(m,$k-p$,p) where p stands for the number of error approximations that are recycled. Typically, p varies between 1 and 3. Note that we want to keep the size k of the recycled space unchanged, irrespective of the small number of recycled error, that is, $p < k$. Baker [19] proposed a numerical criterion based on both sequential and skip angles to analyze the co-linearity of residuals between cycles. We recall a simple expression of the sequential and skip angles relying on the residual norm (see [18]). Let r_{s+1} and r_s be the (s+1)st and the (s)th residual vectors from LGCRO-DR(m,$k-p$,p). Then the sequential angle between two residual vectors is

$$cos\angle(r_{s+1}, r_s) = \frac{\|r_{s+1}\|}{\|r_s\|}. \tag{35}$$

Also, let r_{s+1} and r_{s-1} be the (s+1)st and the (s-1)st residual vectors from LGCRO-DR(m,$k-p$,p). Then the skip angle between two residual vectors is

$$cos\angle(r_{s+1}, r_{s-1}) = \frac{\|r_{s+1}\|}{\|r_{s-1}\|}. \tag{36}$$

We point out that the expression of the skip angle (36) is only valid for the LGCRO-DR method. Otherwise, the skip angle is formulated for the non-flexible GCRO-DR as:

$$cos\angle(r_{s+1}, r_{s-1}) = \frac{\|r_{s+1}\|}{\|r_{s-1}\|} - \frac{< Az_{s+1}, Az_s >}{\|r_{s+1}\|\|r_{s-1}\|}, \tag{37}$$

where $z_{s+1} = x_{s+1} - x_s \in \mathcal{K}_m(A, r_s)$ stands for the error approximation of the (s+1)st cycle.

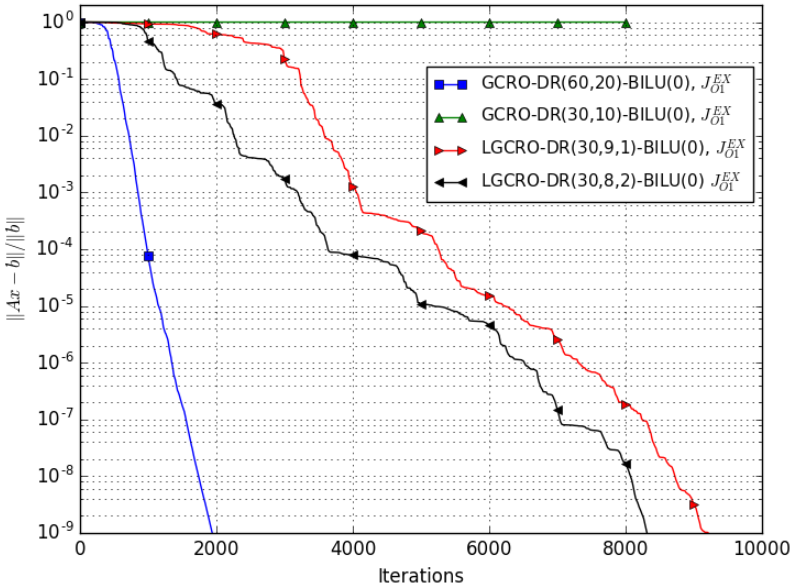

Fig. 13 Adjoint relative residual norm convergence history for GCRO-DR and LGCRO-DR. Impact of recycling the error approximations on the convergence. No scaling strategy is considered in this case.

Fig. 13 shows how the recycling of the error approximation improves the robustness of the GCRO-DR solver. In order to highlight the interest of the error approximation strategy, we decided to solve the aerodynamic adjoint system for two different sizes of the Krylov space. We consider both GCRO-DR(60,20) and GCRO-DR(30,10). As we can see in Fig. 13, GCRO-DR(60,20) converges after 1900 iterations. Even though a plateau is observed at the beginning, the convergence of GCRO-DR(60,20) seems to follow the superlinear convergence property of GMRES without restarting. Fig. 14a depicts the variation of the sequential and skip angles for GCRO-DR(60,20) throughout the cycles. At the beginning, we observe for both sequential and skip angles low values which corroborates the plateau that we see in Fig. 13.

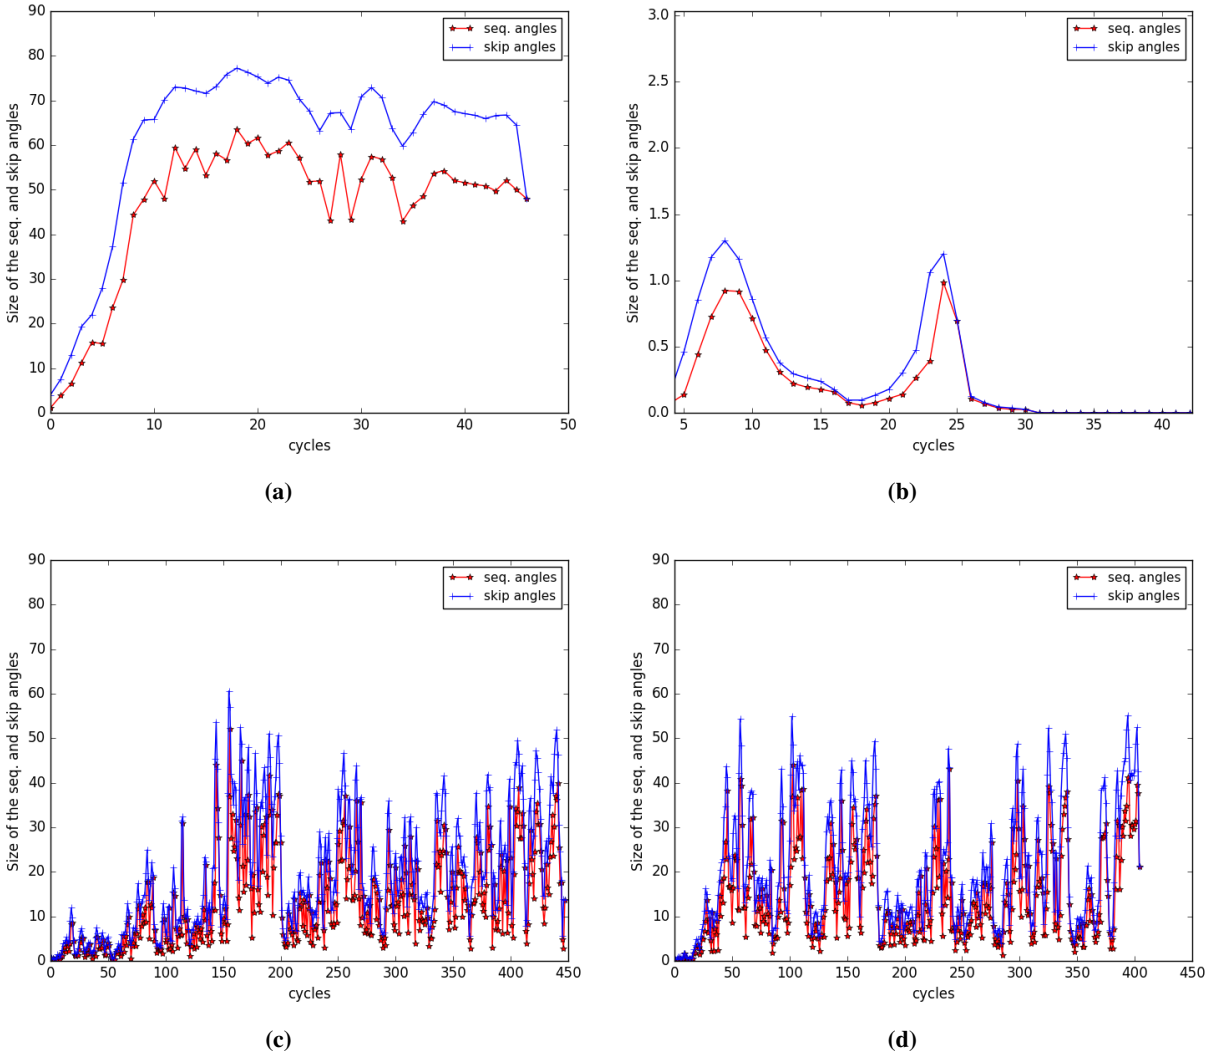

Fig. 14 Sequential and skip angles: (a) GCRO-DR(60,20), (b) GCRO-DR(30,10), (c) LGCRO-DR(30,9,1) and (d) LGCRO-DR(30,8,2).

In contrast, sequential and skip angles are gradually improved throughout the cycles. More precisely, the skip angle reaches up to 80 degrees whereas the sequential angle reaches up to 60 degrees. The second test case consists in reducing the size of the Krylov space by considering GCRO-DR(30,10). In that case, the GCRO-DR solver stagnates. Indeed, this stagnation is clearly illustrated in Fig.14b where the sequential and skip angles never exceed 2 degrees. Let see how the recycling of error approximations affects the convergence behaviour. More specifically, a maximum of two error approximations have been recycled. The red curve in Fig. 13 shows numerically the capability of one error approximation to retrieve the convergence after 9000 iterations. We point out that we have computed the error

approximation in an economical way (see [18]). This economical method has the advantage of being straightforward and easy to implement but is inexact. The LGCRO-DR(30,8,2) solver outperforms the LGCRO-DR(30,9,1) and improves even more the solution with a gain of 10 %. Also, Fig. 14c and 14d shows the correlation between the convergence and the value of the angles. At first glance, there is no major difference between the two Figures. However, if we look a little closer, we see that both angle values in Fig. 14d are higher than those in Fig. 14c. Practically speaking, the lower the angle values the larger the plateau at the beginning of the cycle.

It reflects exactly the size of the plateau of red and black curves in the Fig. 13. Apart from that, no difference appears.

VI. Conclusion

In this exploratory work, recycling Krylov subspace methods have been investigated for improving the partitioned solution of aerostructural adjoint systems.

Numerical experiments have been conducted on turbulent transonic flows over the three-dimensional ONERA M6 wing. In order to challenge our advanced Krylov solvers, a very flexible structural model of the wing was designed leading to a strongly coupled aeroelastic problem. Many of the improvements investigated in this study have proven beneficial to the convergence of Krylov solvers but often for problems of limited size or complexity.

At the beginning of this study, the current strategy was based on a FGMRES-DR solver for the fluid block with a varying right-hand side, i.e., the contribution of the structural source terms is updated at each fluid-structure coupling. However, this class of Krylov solvers is not able to take advantage of valuable spectral information produced during the previous fluid-structure cycles. This is mainly due to the fact that the deflation strategy is based on subspace augmentation. We then decided to implement a projection-based deflated Krylov solver which is able to virtually recycle any subspace at hand. The The current work particularly focused on the FGCRO-DR solver. More specifically, this class of solvers can recycle a spectral information from a linear system to another one.

A first step consisted in developing relevant numerical ingredients in order to improve robustness and efficiency of our inner-outer Krylov solvers. The combination of varying preconditioning strategies associated to proper diagonal scaling of the system matrix led to substantial improvement in the convergence rate, especially when the BILU(0) approximate factorization is applied to the exact first-order flux Jacobian matrix.

These numerical ingredients have been introduced in FGMRES-DR and FGCRO-DR and a numerical comparison showed the equivalence of these two solvers in terms of convergence rate and CPU cost. One important aspect which deserves to be mentioned is that scaling of the system matrix dramatically reduces the plateau of the residual convergence curve, which is a desirable feature in the framework of partitioned solvers when we want to achieve fast convergence between two fluid-structure couplings.

In the context of FGCRO-DR, we then explored subspace recycling from previous fluid-structure cycles. This showed very satisfactory convergence improvements, providing that the recycling process is activated after sufficient convergence of the fixed-point solver. Indeed, effective recycling of spectral information for the solution of a sequence of linear systems assumes that system matrix and right-hand side vary slowly. We proposed some distance criterion to monitor the beginning of recycling. We demonstrated a beneficial impact of recycling on the coupled adjoint system with gain of 20 % in terms of number of iterations.

Finally, we also investigated the recycling of error approximations for the aerodynamic adjoint system and for the standard (non flexible) GCRO-DR solver. Sequential and skip angles have been computed as a relevant criteria for appreciating the convergence. This loose formulation also brought significant improvements to the convergence rate of the Krylov solver. The next step will be to extend the error approximation recycling theory to the flexible class of GCRO solver and to demonstrate its added value in the contex of our partitioned coupled-adjoint solver.

Appendix A Flexible GCRO with deflated restarting (FGCRO-DR) for a sequence of linear systems with varying right-hand sides.

Algorithm 2 FGCRO-DR(m,k)

Choose m, the maximum size of the subspace, and k, the desired number of approximate eigenvectors. let *tol* be the convergence tolerance. Choose an initial guess x_0. Compute $r_0 = b - Ax_0$, and set $i = 1$.
if C_k and Z_k are defined (from solving a previous linear system) **then**
  $x_1^{(i)} = x_0^{(i)} + Z_k C_k^T r_0^{(i)}$ ▷ Do not recompute Z_k and C_k because only b varies
  $r_1^{(i)} = r_0^{(i)} - C_k C_k^T r_0^{(i)}$
else
  $v_1 = r_0 / \|r_0\|_2$
  $c = \|r_0\|_2 e_1$
  Perform m steps of FGMRES, solving $\min \|c - \bar{H}_m y\|_2$ for y and generating V_{m+1}, Z_m and $\bar{H}_m$.
  $x_1 = x_0 + Z_m y$
  $r_1 = V_{m+1}(c - \bar{H}_m y)$
  Set $W_m = V_m$
  Compute the k eigenvectors $\tilde{g}_j$ of $(H_m + h_{m+1,m}^2 H_m^{-T} e_m e_m^T)\tilde{g}_j = \tilde{\theta}_j \tilde{g}_j$ associated with the smallest magnitude eigenvalues $\tilde{\theta}_j$ and store in P_k.
  Let [Q,R] be the reduced QR-factorization of $\bar{H}_m P_k$.
  $C_k = V_{m+1} Q$
  $Z_k = Z_m P_k R^{-1}$
  $W_k = W_m P_k R^{-1}$
end if
while $\|r_i\|_2 > tol$ **do**
  $i = i + 1$
  Perform $(m - k)$ Arnoldi steps with the linear operator $(I - C_k C_k^T)A$ and $v_{k+1} = r_{i-1}/\|r_{i-1}\|_2$ generating V_{m-k+1}, Z_{m-k}, $\bar{H}_{m-k}$, and B_{m-k}.
  $V_{m+1} = [C_k, V_{m-k+1}]$
  $Z_m = [Z_k, Z_{m-k}]$
  $W_m = [W_k, V_{m-k}]$
  $\bar{G}_m = \begin{bmatrix} I_k & B_{m-k} \\ 0 & \bar{H}_{m-k} \end{bmatrix}$
  Solve $\min \|V_{m+1}^T r_{i-1} - \bar{G}_m y\|_2$ for y.
  $x_i = x_{i-1} + Z_m y$
  $r_i = b - Ax_i$
  Compute the k eigenvectors $\tilde{g}_j$ of $\bar{G}_m^T \bar{G}_m \tilde{g}_j = \tilde{\theta}_j \bar{G}_m^T V_{m+1}^T W_m \tilde{g}_j$ associated with smallest magnitude eigenvalues $\tilde{\theta}_j$ and store in P_k.
  Let [Q,R] be the reduced QR factorization of $\bar{H}_m P_k$.
  $C_k = V_{m+1} Q$
  $Z_k = Z_m P_k R^{-1}$
  $W_k = W_m P_k R^{-1}$
end while
Recycling of Z_k and C_k (for the next system)

References

[1] Zhang, Z., and Zingg, D., "Efficient monolithic solution algorithm for high-fidelity aerostructural analysis and optimization," *AIAA Journal*, Vol. 56, No. 3, 2018, pp. 1251–1265.

[2] De Sturler, E., "Truncation strategies for optimal Krylov subspace methods," *SIAM Journal on Numerical Analysis*, Vol. 36, No. 3, 1999, pp. 864–889.

[3] Kenway, G., Kennedy, G., and Martins, J., "Scalable parallel approach for high-fidelity steady-state aeroelastic analysis and adjoint derivative computations," *AIAA journal*, Vol. 52, No. 5, 2014, pp. 935–951.

[4] Saad, Y., and Schultz, M., "GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems," *SIAM Journal on scientific and statistical computing*, Vol. 7, No. 3, 1986, pp. 856–869.

[5] Saad, Y., "A flexible inner-outer preconditioned GMRES algorithm," *SIAM Journal on Scientific Computing*, Vol. 14, No. 2, 1993, pp. 461–469.

[6] Vuik, C., Segal, A., and Meijerink, J., "An efficient preconditioned CG method for the solution of a class of layered problems with extreme contrasts in the coefficients," *Journal of Computational Physics*, Vol. 152, No. 1, 1999, pp. 385–403.

[7] Tang, J., and Vuik, C., "On deflation and singular symmetric positive semi-definite matrices," *Journal of computational and applied mathematics*, Vol. 206, No. 2, 2007, pp. 603–614.

[8] Gaul, A., "Recycling Krylov subspace methods for sequences of linear systems : Analysis and applications," Doctoral thesis, Technische Universität Berlin, Fakultät II - Mathematik und Naturwissenschaften, Berlin, 2014. https://doi.org/10.14279/depositonce-4147, URL http://dx.doi.org/10.14279/depositonce-4147.

[9] Morgan, R., "GMRES with deflated restarting," *SIAM Journal on Scientific Computing*, Vol. 24, No. 1, 2002, pp. 20–37.

[10] Coulaud, O., Giraud, L., Ramet, P., and Vasseur, X., "Deflation and augmentation techniques in Krylov linear solvers," Research Report RR-8265, Feb. 2013. URL https://hal.inria.fr/hal-00803225, preliminary version of the book chapter entitled "Deflation and augmentation techniques in Krylov linear solvers" published in "Developments in Parallel, Distributed, Grid and Cloud Computing for Engineering", ed. Topping, B.H.V and Ivanyi, P., Saxe-Coburg Publications, Kippen, Stirlingshire, United Kingdom, ISBN 978-1-874672-62-3, p. 249-275, 2013.

[11] Giraud, L., Gratton, S., Pinel, X., and Vasseur, X., "Flexible GMRES with deflated restarting," *SIAM Journal on Scientific Computing*, Vol. 32, No. 4, 2010, pp. 1858–1878.

[12] Simoncini, V., and Szyld, D. B., "Flexible inner-outer Krylov subspace methods," *SIAM Journal on Numerical Analysis*, Vol. 40, No. 6, 2002, pp. 2219–2239.

[13] Jadoui, M., Blondeau, C., Martin, E., Renac, F., and Roux, F.-X., "Comparative Study of Inner-Outer Krylov Solvers for Linear Systems in Structured and High-Order Unstructured CFD Problems," *14th WCCM-ECCOMAS Congress 2020*, Vol. 700, 2021.

[14] Parks, M., De Sturler, E., Mackey, G., Johnson, D., and Maiti, S., "Recycling Krylov subspaces for sequences of linear systems," *SIAM Journal on Scientific Computing*, Vol. 28, No. 5, 2006, pp. 1651–1674.

[15] de Sturler, E., "Nested Krylov methods based on GCR," *Journal of Computational and Applied Mathematics*, Vol. 67, No. 1, 1996, pp. 15–41.

[16] Van der Vorst, H., and Vuik, C., "GMRESR: a family of nested GMRES methods," *Numerical linear algebra with applications*, Vol. 1, No. 4, 1994, pp. 369–386.

[17] Carvalho, L., Gratton, S., Lago, R., and Vasseur, X., "A flexible generalized conjugate residual method with inner orthogonalization and deflated restarting," *SIAM Journal on Matrix Analysis and Applications*, Vol. 32, No. 4, 2011, pp. 1212–1235.

[18] Niu, Q., Lu, L.-Z., and Liu, G., "Accelerated GCRO-DR Method for Solving Sequences of Systems of Linear Equations," *Journal of computational and applied mathematics*, Vol. 253, 2013, pp. 131–141.

[19] Baker, A., Jessup, E., and Manteuffel, T., "A technique for accelerating the convergence of restarted GMRES," *SIAM Journal on Matrix Analysis and Applications*, Vol. 26, No. 4, 2005, pp. 962–984.

[20] Achard, T., "Techniques de calcul de gradient aéro-structure haute-fidélité pour l'optimisation des voilures flexibles," Ph.D. thesis, ONERA/CNAM LMSSC, Dec. 2017. URL http://www.theses.fr/2017CNAM1140.

[21] Achard, T., Blondeau, C., and Ohayon, R., "High-Fidelity Aerostructural Gradient Computation Techniques with Application to a Realistic Wing Sizing," *AIAA Journal*, Vol. 56, No. 11, 2018, pp. 4487–4499.

[22] Axelsson, O., *Iterative solution methods*, Cambridge university press, 1996.

[23] Eisenstat, S., Elman, H., and Schultz, M., "Variational iterative methods for nonsymmetric systems of linear equations," *SIAM Journal on Numerical Analysis*, Vol. 20, No. 2, 1983, pp. 345–357.

[24] Cambier, L., Heib, S., and Plot, S., "The Onera elsA CFD software: input from research and feedback from industry," *Mechanics & Industry*, Vol. 14, No. 3, 2013, pp. 159–174. https://doi.org/10.1051/meca/2013056, URL https://hal.archives-ouvertes.fr/hal-01293795.

[25] Cambier, L., Gleize, V., and Julien, M., *Verification and validation of the Onera elsA flow solver on RANS benchmarks*, ???? https://doi.org/10.2514/6.2014-0239, URL https://arc.aiaa.org/doi/abs/10.2514/6.2014-0239.

[26] Saad, Y., *Iterative Methods for Sparse Linear Systems*, SIAM, 2003. https://doi.org/10.1137/1.9780898718003.

[27] Peter, J., and Drullion, F., "Large stencil viscous flux linearization for the simulation of 3D compressible turbulent flows with backward-Euler schemes," *Computers & Fluids*, Vol. 36, No. 6, 2007, pp. 1005–1027.

[28] Cai, X.-C., and Sarkis, M., "A Restricted Additive Schwarz Preconditioner for General Sparse Linear Systems," *SIAM J. Sci. Comput.*, Vol. 21, No. 2, 1999, pp. 792–797. https://doi.org/10.1137/s106482759732678x.

[29] Yoon, S., and Jameson, A., "Lower-upper symmetric-Gauss-Seidel method for the Euler and Navier-Stokes equations," *AIAA journal*, Vol. 26, No. 9, 1988, pp. 1025–1026.