

COMPARATIVE STUDY OF INNER-OUTER KRYLOV SOLVERS FOR LINEAR SYSTEMS IN STRUCTURED AND HIGH-ORDER UNSTRUCTURED CFD PROBLEMS

M. JADOU^{1*}, C. BLONDEAU^{1†}, E. MARTIN^{1◦}, F. RENAC^{1‡} and F.-X. ROUX^{1⊙}

¹ ONERA, Université Paris-Saclay, F-92322 Châtillon, www.onera.fr

*Mehdi.Jadoui@onera.fr
†Christophe.Blondeau@onera.fr
◦Emeric.Martin@onera.fr
‡Florent.Renac@onera.fr
⊙François-Xavier.Roux@onera.fr

Key words: Krylov solvers, GMRES, adjoint problem, deflation

Abstract. Advanced Krylov subspace methods are investigated for the solution of linear systems arising from an adjoint-based aerodynamic shape optimization problem. A special attention is paid for the flexible inner-outer GMRES strategy combined with most relevant preconditioning strategies and deflation techniques. The choice of this specific class of Krylov solvers for solving challenging problems is based on its outstanding convergence properties. Moreover, parallel scalability is improved by globalizing the preconditioning phase through an additive domain decomposition technique. However, maintaining the performance of the preconditioner may be challenging since scalability and efficiency of a preconditioning technique are properties often antagonistic to each other. We demonstrate how flexible inner-outer Krylov methods are able to overcome this critical issue. A numerical comparative study is provided on the supercritical OAT15A airfoil in turbulent flow under transonic regime conditions using a Finite Volume method (FV) and a High-Order Discontinuous Galerkin (DG) one. Based on this representative problem a discussion of the recommended numerical practices is proposed.

1 INTRODUCTION

The solution of large sparse linear systems is a substantive subject in many large-scale simulations especially in computational fluid dynamics (CFD). Due to their prohibitive memory requirement, direct methods are often overlooked in favor of iterative methods when linear systems become very large and a wide range of iterative linear solvers exists today to tackle such kind of problems. The most popular are the preconditioned Krylov subspace Method [1]. For non symmetric problems, the Generalized Minimal Residual method GMRES [2] and its flexible variant FGMRES [3] are preferred due to their robustness and their expected superlinear convergence property. Unfortunately, the GMRES algorithm requires significant memory due to the storage of a sequence of orthogonal vectors to approximate the solution. To limit

the number of basis vectors, we proceed by a restarted version of the algorithm referred to as GMRES(m), where m is the maximum number of basis vectors basis allowed during a cycle. Once m iterations are performed, the vectors are discarded and a new set of vectors is generated. Cycles are repeated until a prescribed tolerance is reached.

An essential numerical ingredient that ensures good convergence properties of Krylov subspace methods is the quality of the preconditioner. Basically, one attempts to transform the initial linear system into another one easier to solve. From a spectral point of view, a good preconditioner ensures as much as possible the clustering of the eigenspectrum of the initial system and provides a good approximation of the inverse of the initial system matrix. Benzi gives a comprehensive overview of preconditioning techniques for the iterative solution of large linear systems [4]. For instance, the Incomplete Lower-Upper (ILU) factorization and its block variant (BILU) were specifically studied. ILU simply consists in the LU factorization of a sparse matrix with dropping strategy. The dropping strategy is based on several criteria, such as the location or magnitude of the matrix element. Another popular preconditioner is the Lower-Upper Symmetric Successive Over Relaxation (LU-SSOR) which is an iterative method combining the LU factorization with a relaxation type SSOR method [5].

However, the restarted preconditioned GMRES(m) still lacks robustness and a stagnation of the residuals may be observed. The limitation stems from the restart process of the GMRES(m). Indeed, when a new cycle starts, all previous information related to the Krylov basis is lost. This leads to a non-optimal search in the subsequent iterations resulting in a stalled convergence. Actually, the rate of convergence during first iterations is dramatically hampered by the smallest eigenvalues of the preconditioned system. The associated spectral information is mainly contained in the preceding Krylov subspace. Thus, instead of discarding the previous Krylov subspace information, Morgan [6] proposed to deflate the smallest eigenvalues by augmenting the next Krylov subspace with the harmonic Ritz vectors that approximate the eigenvectors associated to the exterior eigenvalues of the preconditioned system. These harmonic Ritz vectors are formulated as a linear combination of Krylov vectors whose coefficients are solutions of a generalized eigenvalue problem involving the upper Hessenberg matrix [7]. The purpose of deflation is to rule out the harmful effect of smallest (or highest in some cases) eigenvalues which tend to slow down the convergence of the Krylov subspace methods. It leads to a new algorithm denoted as the GMRES-DR(m,k) (GMRES with Deflated Restarting) with k the number of selected harmonic Ritz vectors per cycle. It was proved to improve the convergence rate of GMRES(m) (see [8] for the Flexible variant FGMRES-DR).

This paper is organized as follows. In section 2, we briefly present the adjoint problem for a sensitivity analysis. In sections 3 and 4, we describe the minimum residual norm Krylov subspace methods in conjunction with deflation strategies and flexible preconditioning. Section 5 is devoted to numerical experiments where the robustness of flexible inner-outer Krylov solvers is demonstrated through a representative test case for both FV ¹ and DG ² schemes. Finally, strong scaling capability of such solvers is assessed.

¹FV simulations were performed with the ONERA elsA software [9]

²DG simulations were performed with ONERA Aghora software [10]

2 PROBLEM STATEMENT

2.1 Steady-state problem

Let \mathcal{D} be a bounded domain in \mathbb{R}^3 . We denote by $\mathbf{W} \in \mathbb{R}^{n_a}$ the vector of conservative variables. The differential form of the governing equations for a viscous fluid governed by the RANS model is formulated as

$$\frac{d\mathbf{W}}{dt} + \nabla \cdot \mathbf{F} = \mathbf{Q}, \quad (1)$$

where \mathbf{F} embeds the convective and diffusive fluxes and \mathbf{Q} is a source term. We want to obtain the field of conservative variables \mathbf{W} related to a steady-state equilibrium on domain \mathcal{D} .

The system of equations is formulated in it's discrete residual form as

$$\mathbf{R}(\mathbf{W}, \mathbf{X}) = 0, \quad (2)$$

where $\mathbf{X} \in \mathbb{R}^N$ is the vector of fluid mesh coordinates.

2.2 Sensitivity analysis via the adjoint approach

Consider an aerodynamic shape optimization problem of the form $\mathcal{J}^* = \min_{\alpha} \mathcal{J}(\mathbf{W}(\alpha), \mathbf{X}(\alpha))$ under the constraints $\mathbf{R}(\mathbf{W}(\alpha), \mathbf{X}(\alpha)) = 0$, where α is the vector of design variables controlling the shape of the body. A gradient-based method is required to efficiently minimize \mathcal{J} . In aerodynamics, the adjoint technique is the method of choice to compute the sensitivities [11].

The adjoint linear system is written as

$$\begin{bmatrix} \partial \mathbf{R} \\ \partial \mathbf{W} \end{bmatrix}^T \lambda = -\frac{\partial \mathcal{J}^T}{\partial \mathbf{W}} \quad (3)$$

In the remainder of this paper, the exact flux Jacobian matrix will be noted as $A = \frac{\partial \mathbf{R}^T}{\partial \mathbf{W}}$.

3 MINIMAL RESIDUAL KRYLOV SUBSPACE METHOD

In this section we focus on a particular minimal residual norm Krylov subspace method for the solution of linear systems with a non-symmetric real coefficient matrix of type

$$Ax = b, \quad A \in \mathbb{R}^{N \times N}; \quad b, x \in \mathbb{R}^N \quad (4)$$

In the remainder of this paper, a right-preconditioned system will be considered so that system (4) becomes

$$AM^{-1}t = b, \tag{5}$$

$$x = M^{-1}t \tag{6}$$

3.1 GMRES algorithm with right preconditioning

The GMRES algorithm is an iterative method [2] which consists in building an approximation of the solution x_m into a search subspace of dimension $m \ll N$ by minimizing the norm of the associated residual vector $r_m = b - Ax_m$.

We call this search space the Krylov subspace and defined it as

$$\mathcal{K}_m(AM^{-1}, r_0) = \text{span}\{r_0, AM^{-1}r_0, \dots, (AM^{-1})^{m-1}r_0\} \tag{7}$$

where x_0 is an initial guess to the solution x .

More precisely, GMRES generates an orthonormal basis of \mathcal{K}_m denoted by $V_m = [v_1, \dots, v_m]$. This basis is computed by an Arnoldi recurrence based on the Modified Gram-Schmidt orthogonalization procedure. The approximated solution is written as $x_m = x_0 + M^{-1}V_my_m$ with $y_m \in \mathbb{R}^m$. The Arnoldi process leads to the fundamental relationship

$$AM^{-1}V_m = V_{m+1}\tilde{H}_m, \tag{8}$$

with $\tilde{H}_m \in \mathbb{R}^{(m+1) \times m}$ an upper Hessenberg matrix and $V_{m+1} \in \mathbb{R}^{N \times (m+1)}$ such that $V_{m+1}^T V_{m+1} = I_{m+1}$, where the superscript T stands for the conjugate transpose of the matrix.

We want to minimize the residual norm $\|r_m\|_2$ at the m^{th} iteration over the Krylov subspace \mathcal{K}_m . By using (8) and the fact that $x_m = x_0 + M^{-1}V_my_m$, the residual minimization problem leads to the following overdetermined least squares problem of dimension $(m+1)$

$$y_m = \min_{y \in \mathbb{R}^m} \|r_0 - AM^{-1}V_my\|_2 = \min_{y \in \mathbb{R}^m} \|\beta e_1 - \tilde{H}_m y\|_2, \tag{9}$$

where $\beta = \|r_0\|_2$ and e_1 the unit normal vector of the canonical Euclidean basis of \mathbb{R}^{m+1} .

This version of GMRES is actually impracticable due to high memory footprint and significant computational cost. To overcome such limitations, we adopt the restarted GMRES version referred to as GMRES(m) where $m \ll N$. If the stopping criterion does not reach a prescribed tolerance at the outcome of the first m iterations, the algorithm is restarted from an updated initial guess $x_0 = x_m$.

In our study, we focus on the Flexible GMRES (FGMRES) approach proposed by Saad [3]. This approach allows variable preconditioning compared to the usual preconditioned GMRES where the preconditioner is stationary. Allowing a variable preconditioner adds the great flexibility of using iterative solvers to obtain the next search direction. By this way, two embedded

linear systems are defined but at the price of twice the memory cost initially planned for the outer system and an extra memory cost for the inner system dedicated to the preconditioning step.

The flexible Arnoldi relationship is written as

$$AZ_m = V_{m+1}\bar{H}_m, \quad Z_m \in \mathbb{R}^{N \times m} \quad (10)$$

and the approximate solution becomes $x_m = x_0 + Z_m y_m$ where y_m minimizes $\|r_0 - AZ_m y\|_2$ over $x_0 + \text{span}\{Z_m\}$. We point out that a storage has to be allocated for both matrices $Z_m = [M_1^{-1}v_1, \dots, M_m^{-1}v_m]$ and V_m .

The Flexible GMRES algorithm is presented in Algorithm 1. We denote by m_i the Krylov subspace size of the GMRES solver devoted to the inner linear system.

Algorithm 1 FGMRES(m, m_i)

- 1: Choose an initial guess x_0 and a convergence threshold ε
 - 2: Compute $r_0 = b - Ax_0$, $\beta = \|r_0\|$ and $v_1 = r_0/\beta$;
 - 3: **for** $j = 1, \dots, m$ **do**
 - 4: $z_j = M_j^{-1}v_j$
 - 5: $w = Az_j$;
 - 6: **for** $i = 1, \dots, j$ **do**
 - 7: $h_{i,j} = v_i^H w$;
 - 8: $w = w - h_{i,j}v_i$;
 - 9: **end for**
 - 10: $h_{j+1,j} = \|w\|$ and $v_{j+1} = w/h_{j+1,j}$;
 - 11: Solve the least-squares problem $\min_y \|\beta e_1 - \bar{H}_j y\|$ for y^* ;
 - 12: Exit if $\|c - \bar{H}_m y^*\|/\|b\| \leq \varepsilon$
 - 13: **end for**
 - 14: Compute $x_m = x_0 + Z_m y^*$ where $Z_m = [z_1, \dots, z_m]$
 - 15: Set $x_0 = x_m$ and go to 2
-

In this paper the nested GMRES strategy is adopted for all numerical experiments. Two preconditioning strategies are considered for the inner GMRES. The first one consists in using a block-diagonal version of a standard LU-SGS iterative solver. LU-SGS is applied on a first order diagonally dominant upwind approximation of the flux Jacobian matrix inspired by [5]. This operator is based on a first order spatial discretization of the convective and of the viscous fluxes using a simplifying thin layer assumption (see [12]). This strategy leads to a very compact stencil for the preconditioning matrix. The second one is an incomplete BILU(k) factorization applied to various exact or approximate flux Jacobian matrices.

More specifically, BILU(k) will be applied either on the first order approximate Jacobian matrix $\mathbf{J}_{O_1}^{APP}$ or on the first exact Jacobian matrix $\mathbf{J}_{O_1}^{EX}$. $\mathbf{J}_{O_1}^{EX}$ is different from $\mathbf{J}_{O_1}^{APP}$ when it comes to memory footprint. More precisely, $\mathbf{J}_{O_1}^{EX}$ has a stencil of 9-points in 2D whereas a stencil of 5-

points is observed for $\mathbf{J}_{O_1}^{APP}$. Consequently, one needs about twice the storage for $\mathbf{J}_{O_1}^{EX}$ compared to $\mathbf{J}_{O_1}^{APP}$ but with a better robustness. For the High-Order formulation, an exact third order Jacobian matrix $\mathbf{J}_{O_3}^{EX}$ is built. Only BILU(0) is considered in the latter case.

We point out that the relevant numerical ingredients that characterize the GMRES algorithm are the matrix product (step 5), the preconditioning step (step 4) and the scalar product (step 7). These algebraic operations are global in conjunction with a domain decomposition method. More precisely, the globalization of the preconditioner (step 4) is achieved with a Restrictive Additive Schwarz method [13]. In addition, the product by the operator A (step 5) is exact. In the end, we get a global and parallel FGMRES(m, m_i).

3.2 Deflation approach

The main drawback of the restarted GMRES(m) is the loss of crucial information during the restarting procedure. This information is essentially contained in the Hessenberg matrix \bar{H}_m which represents the restriction of the operator AM^{-1} to the Krylov subspace \mathcal{K}_m . Indeed, thanks to (8), we deduce easily that:

$$H_m = V_m^t A M^{-1} V_m \quad (11)$$

Therefore, the spectrum of H_m naturally approximates a part of the spectrum of AM^{-1} , especially the extreme eigenvalues, which are responsible for stagnation. The idea of *deflation* is then to eliminate the unwanted effect of those eigenvalues of AM^{-1} to improve the convergence. In [8], Giraud et al. performed deflation for both smallest and largest eigenvalues to maximize the deflation effect. In contrast, Morgan [6] focused on the smallest ones. The first strategy will be considered in our numerical experiments.

There are two types of deflation. Deflation by projection [14] consists in building a projector P such that the deflated system $P Ax = P b$ is easier to solve than the original system (5). The second type consists in adding a certain number k of harmonic Ritz vectors into the Krylov subspace [6]. They are supposed to provide a good approximation of eigenvectors associated to the smallest eigenvalues of AM^{-1} . For implementation reasons, the second option is investigated and the subsequent solver is denoted by FGMRES-DR(m, m_i, k) (see [8] for implementation details).

4 NUMERICAL RESULTS

To evaluate this GMRES solver capability, a dedicated test case has been defined. We consider a 2D transonic supercritical OAT15A airfoil. The Mach number is 0.734, Reynolds number is 6.5M and the angle of attack is 1.15° . The fluid model is RANS and the Spalart Allmaras one equation turbulence model is selected. We focus on the resolution of the adjoint system with the fully linearized turbulence model. For the FV scheme, the numerical scheme is a second order Roe spatial discretization associated to a MUSCL reconstruction and a Van Albada limiter. With regard to preconditioning step, $\mathbf{J}_{O_1}^{EX}$ and $\mathbf{J}_{O_1}^{APP}$ will be tested. For the DG scheme, the numerical scheme relies on the Roe flux for the convective term and on the Bassi-Rebay-2

(BR2) [15] method for the diffusive term, and $\mathbf{J}_{O_3}^{EX}$ is exploited to build the preconditioner.

4.1 OAT15A airfoil - steady state analysis

The OAT15A profile has a thick trailing edge such that the initial structured mesh is made of two blocks. The structured mesh for the FV calculation is composed of 158 208 hexahedra and the unstructured mesh for the DG one of 64 416 hexahedra.




Figure 1: 2D transonic turbulent flow over the OAT15A airfoil : (a) and (b) represent respectively the 16-block structured and unstructured mesh while the steady state density field convergence is illustrated in (c) and (d) for both FV and DG schemes.


Figure 1a and Figure 1b depict the mesh partitioning obtained on 16 cores for the scalability study. For the unstructured case, a line construction algorithm is first applied to detect strength of couplings between unknowns from a directional stiffness measurement [16]. Then, a multilevel k-way graph partitioning algorithm from the METIS library is applied based on a strategy taking into account weights on both the vertices and the edges of the graph. Finally, four weights have been defined and calibrated to distinguish mesh cells belonging to a strong connection line but also mesh edges whose one of the parent cell is affected. The main interest is to preserve as much as possible certain critical areas of the original domain from being arbitrarily partitioned. For example in the case of a shock wave propagation or a turbulent boundary layer development. The resulting largest load imbalance for the unstructured case finally is of the same order as that obtained for the structured case and is about 8 %.

4.2 OAT15A - Adjoint sensitivity analysis, fully linearized turbulence model

Recall that we want to solve the adjoint system (3) efficiently by using the inner-outer GMRES solver with deflated restarting denoted by FGMRES-DR(m, m_i, k). We adopt the relevant numerical parameters in Table 1 below:

m	60	Size of the outer Krylov space
k	20	Number of deflated vectors
tol outer	1.e-9	Relative convergence threshold
m_i	20	Maximum size of the inner Krylov space
tol inner	0.5	Convergence threshold for the inner GMRES

Table 1: FGMRES-DR relevant numerical parameters



(a) Size of the FV problem $N \sim 1M$ with $NNZ \sim 35M$ (b) Size of the DG problem $N \sim 2M$ with $NNZ \sim 288M$

Figure 2: Adjoint density residual norm convergence history, fully linearized SA turbulence model.

All computations are performed with 16 cores on 16-blocks structured and unstructured grid thus with one fluid domain associated to each core. Figure 2 displays the residual norm convergence history of FGMRES-DR for both FV and DG schemes. In the FV case, we notice a stagnation of FGMRES(60,20) for both LU-SGS and BILU(0) algorithms applied to the first order approximate Jacobian matrix J_{O1}^{APP} . The natural choice to overcome this stagnation is to extend the projection space \mathcal{K}_m in order to capture as much as possible the eigenvalues of AM^{-1} . Since we cannot afford to extend the projection space due to memory limitation, the deflation strategy remains the best alternative to retrieve the residual convergence of a quasi-full FGMRES solver at low computational cost. A deflation of 30 % of m is sufficient and the resulting FGMRES-DR(60,20,20) solver, preconditioned by LU-SGS(6) or BILU(0) choice now converges as we can see in the Figure 2a. However, when it comes to use the first order exact Jacobian matrix J_{O1}^{EX} , FGMRES(60,20) converges in a few iterations. In this case, the de-

flation does not lead to any improvements compared to FGMRES(60,20) which illustrates how robust \mathbf{J}_{O1}^{EX} is. The shortcoming of this is that we need approximately twice the storage of \mathbf{J}_{O1}^{APP} . Figure 2b shows convergence history of both FGMRES(60,20,20) and FGMRES-DR(60,20,20) for the DG case. From a DG formulation, resulting Jacobian matrices are often ill-conditioned, especially when the spatial order of the scheme becomes high. A combination of the BILU(0) algorithm with a Restrictive Additive Schwarz method has been done in order to globalize as much as possible the preconditioning effect on the whole fluid domain (BILU(0)+RAS). We have reported the performance details of the two solvers in Table 2 and Table 3 for both FV and DG formulations.

$M^{-1} = \text{BILU}(0)$	FGMRES(60,20)		FGMRES-DR(60,20,20)	
Jacobian matrix	\mathbf{J}_{O1}^{EX}	\mathbf{J}_{O1}^{APP}	\mathbf{J}_{O1}^{EX}	\mathbf{J}_{O1}^{APP}
memory (Go)	0.26	0.21	0.26	0.21
# its	72	-	83	546

Table 2: Performance of FGMRES-DR for FV

$M^{-1} = \text{BILU}(0)+\text{RAS}$	FGMRES(60,20)	FGMRES-DR(60,20,20)
Jacobian matrix	\mathbf{J}_{O3}^{EX}	\mathbf{J}_{O3}^{EX}
memory (Go)	2.1	2.1
# its	111	93

Table 3: Performance of FGMRES-DR for DG

It is worth noticing that for FV case the matrix-vector product involving the Jacobian matrix is performed on the fly. Therefore, only the memory cost required to build the preconditioner is reported in both Table 2 and Table 3.

4.3 Scalability

We have demonstrated in the previous section how competitive the inner-outer GMRES with deflation techniques are for stiff problems. However, it is very common to measure the scalability of the solver in a parallel environment. Here, we pay a special attention to the strong scalability in the sense that the problem size is fixed while varying the number of cores.

For the FV scheme, we notice on Figure 3a that the speedup reaches up to 75 % for the LU-SGS(6) applied to \mathbf{J}_{O1}^{APP} while the speedup for the BILU(0) applied to \mathbf{J}_{O1}^{EX} reaches up to 80 %. The same results have been observed for the DG scheme since a speedup of 85 % for FGMRES(60,20) as well as for the FGMRES-DR(60,20,20) is measured (Figure 3b). We point out that the local nature of the BILU(0) preconditionner does not have a real impact on the robustness of the inner-outer GMRES methods unlike the standard GMRES [17]. Indeed, one of the main advantages of the inner-outer GMRES is the global nature of the inner GMRES preconditionner since it solves the same initial linear system (4).

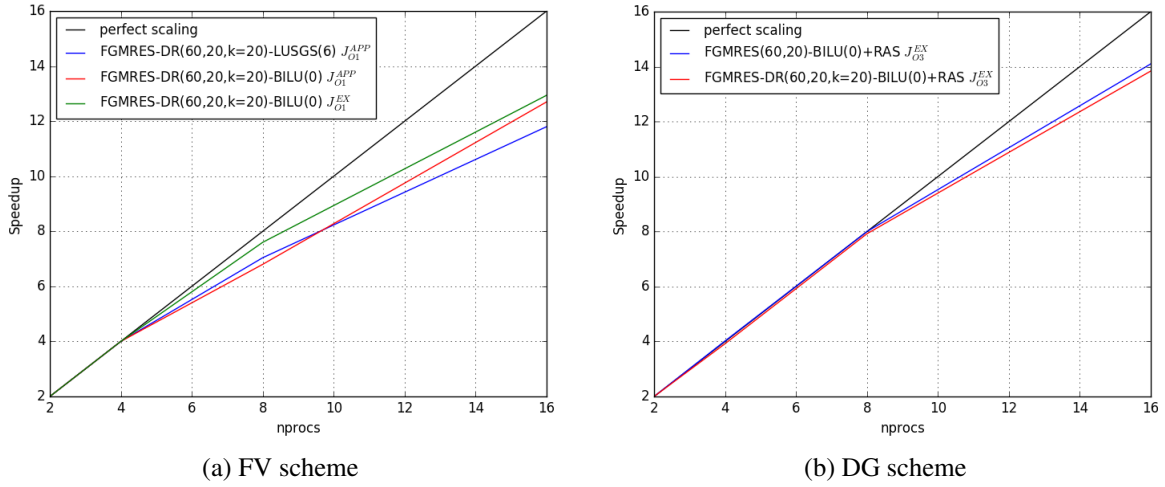


Figure 3: Strong scalability analysis for FGMRES-DR(60,20,20) solvers on OAT15A adjoint system.

4.4 Convergence analysis of inner-outer Krylov methods

It's worthwhile to highlight the convergence property of the inner-outer GMRES algorithm compared to the restarted version. It is shown in [17] that the subspace from where the approximations are chosen keeps growing as the outer iteration j progresses. Indeed, as noted earlier in section 3.1, the global solution x_j is a linear combination of the columns of Z_j . Since each z_j approximates the solution of (11) using an inner GMRES(m_j), a Krylov subspace $\mathcal{K}_{m_j}(A, w_j)$ where w_j stands for the j^{th} outer Krylov vector, is generated with its corresponding basis $[v_1^{(j)}, \dots, v_{m_j}^{(j)}]$. The matrix $V_m^{(j)} = [v_1^{(j)}, \dots, v_{m_j}^{(j)}]$ satisfies the Arnoldi recurrence (7), i.e,

$$AV_{m_j}^{(j)} = V_{m_j+1}^{(j)} H_{m_j}^{(j)}, \quad (12)$$

$$z_j = V_{m_j}^{(j)} y_j, \quad y_j \in \mathbb{R}^{m_j} \quad (13)$$

From the fact that $x_j = x_0 + Z_j u_j$ for some $u_j \in \mathbb{R}^j$ and the relation (13), we conclude that the global solution x_j can be expressed as a linear combination of all the bases of the j inner Krylov subspaces. If we collect all the j inner basis into a global matrix $\mathcal{B}_j \in \mathbb{R}^{N \times p}$, it follows that

$$x_j = x_0 + \mathcal{B}_j Y_j u_j, \quad (14)$$

where p is given by $p = \sum_{i=1}^j m_i$ and $Y_j = \text{diag}(y_1, \dots, y_j) \in \mathbb{R}^{p \times j}$. It should be noted that the columns of \mathcal{B}_j are not necessarily linearly independent.

5 DISCUSSION

In our numerical experiments we have considered tiny convergence thresholds because it is acknowledged that adjoint solutions may exhibit oscillations or still be locally under-resolved

for moderate relative residual decrease (typically 3 orders or magnitude). This strong convergence is only attainable in practice if a number of precautions was taken in the implementation of the FGMRES-DR solver. First, the orthogonality of the Krylov basis at the outcome of the Arnoldi process or after a deflation step must be preserved as much as possible to ensure convergence in a reasonable number of iterations for stiff problems. This point was studied in [18] where the authors advocate the use of a re-orthogonalization during the MGS process. Morgan [6] also proposed an economical strategy where only the last Krylov vector of the deflated basis is re-orthogonalized with respect to other deflated basis vectors. The initial physical equations should always be solved in their non dimensionalized form. For challenging matrices Saad [19] also recommends to apply a scaling to all the rows (or columns) e.g., so that their 1-norms are all equal to 1; then to apply a scaling of the columns (or rows). This practice is strongly recommended before any ILU type factorization. Additionally, recognizing that the linearization of the turbulence model dramatically deteriorates the condition number of the system matrix, Chisholm and Zingg [20] advice to apply a specific scaling coefficient to this particular equation. Additional numerical ingredients like the mixed-precision nested GMRES, i.e. single precision floating point arithmetics for the inner GMRES, and the application of the BILU factorization to the RAS partition of the fluid domain have proven beneficial. Last but not least, the availability of different physical approximations to build the candidate flux Jacobian matrix for the preconditioner is very interesting to keep a good compromise between efficiency and memory footprint as it has a strong impact on the bandwidth of the stencil of the operator. Finally, some memory saving improvements were proposed in [9] to perform in place the matrix-matrix product leading to the new deflated Krylov vector basis at the starting of the next cycle. In fact in a nested FGMRES-DR framework this is not required anymore because the temporary memory required by this product is already available for the storage of the inner Krylov space and is then naturally re-used during the outer level deflation step. It is also worth mentioning that in practice the inner GMRES solver is never restarted considering the small size of the inner Krylov subspace.

6 CONCLUSION

An efficient inner-outer Krylov subspace method with deflated restarting has been proposed for solving adjoint problems when the standard GMRES fails. Relevant numerical tools have been also developed for improving the efficiency of the FGMRES as the deflation strategy and preconditioning techniques. For the latter one, we have also investigated the impact of using different physical approximations of the flux Jacobian matrix for the preconditioning step. Performance of such an approach has been evaluated on a 2D transonic supercritical OAT15A airfoil using Finite Volume and Discontinuous Galerkin schemes. The efficiency of the nested GMRES has been observed when deflation technique with the first order approximate Jacobian matrix is considered or by using the first order exact Jacobian matrix without deflation but at the price of twice the memory cost. Also, we have exhibited a good parallel scalability of the FGMRES thanks to the inherent global effect of the inner GMRES. The interesting conclusion of this work is the capability of inner-outer GMRES strategy to solve stiff problems and exhibit the superlinear convergence of the Full GMRES with a restricted memory footprint.

References

- [1] V. Simoncini and D.B. Szyld. “Recent computational developments in Krylov subspace methods for linear systems”. In: *Numerical Linear Algebra with Applications* 14.1 (2007), pp. 1–59.
- [2] Y. Saad and M.H. Schultz. “GMRES: A generalized minimal residual algorithm for solving non-symmetric linear systems”. In: *SIAM Journal on scientific and statistical computing* 7.3 (1986), pp. 856–869.
- [3] Y. Saad. “A flexible inner-outer preconditioned GMRES algorithm”. In: *SIAM Journal on Scientific Computing* 14.2 (1993), pp. 461–469.
- [4] M. Benzi. “Preconditioning techniques for large linear systems: a survey”. In: *Journal of Computational Physics* 182.2 (2002), pp. 418–477.
- [5] S. Yoon and A. Jameson. “Lower-upper symmetric-Gauss-Seidel method for the Euler and Navier-Stokes equations”. In: *AIAA journal* 26.9 (1988), pp. 1025–1026.
- [6] R.B. Morgan. “GMRES with deflated restarting”. In: *SIAM Journal on Scientific Computing* 24.1 (2002), pp. 20–37.
- [7] R.B. Morgan and M. Zeng. “Harmonic projection methods for large non-symmetric eigenvalue problems”. In: *Numerical linear algebra with applications* 5.1 (1998), pp. 33–55.
- [8] L. Giraud et al. “Flexible GMRES with deflated restarting”. In: *SIAM Journal on Scientific Computing* 32.4 (2010), pp. 1858–1878.
- [9] L. Cambier, S. Heib, and S. Plot. “The Onera elsA CFD software: input from research and feedback from industry”. In: *Mechanics & Industry* 14.3 (2013), pp. 159–174.
- [10] F. Renac et al. “Aghora: A high-order DG solver for turbulent flow simulations”. In: 128 (2015).
- [11] A. Jameson. “Optimum aerodynamic design using CFD and control theory”. In: *12th computational fluid dynamics conference*. 1995, p. 1729.
- [12] J. Peter and F. Drullion. “Large stencil viscous flux linearization for the simulation of 3D compressible turbulent flows with backward-Euler schemes”. In: *Computers & Fluids* 36.6 (2007), pp. 1005–1027.
- [13] X.-C. Cai and M. Sarkis. “A Restricted Additive Schwarz Preconditioner for General Sparse Linear Systems”. In: *SIAM Journal on Scientific Computing* 21.2 (1999), pp. 792–797.
- [14] J. Erhel, K. Burrage, and B. Pohl. “Restarted GMRES preconditioned by deflation”. In: *Journal of computational and applied mathematics* 69.2 (1996), pp. 303–318.
- [15] F. Bassi et al. “A high-order accurate discontinuous finite element method for inviscid and viscous turbomachinery flows”. In: *Proceedings of the 2nd European Conference on Turbomachinery Fluid Dynamics and Thermodynamics*. Antwerpen, Belgium. 1997, pp. 99–109.
- [16] B.R. Ahrabi and D. Mavriplis. “An implicit block ILU smoother for preconditioning of Newton-Krylov solvers with application in high-order stabilized finite-element methods”. In: *Computer Methods in Applied Mechanics and Engineering* 358 (2020).
- [17] V. Simoncini and D.B. Szyld. “Flexible inner-outer Krylov subspace methods”. In: *SIAM Journal on Numerical Analysis* 40.6 (2002), pp. 2219–2239.
- [18] S. Röllin and W. Fichtner. “Improving the accuracy of GMRES with deflated restarting”. In: *SIAM Journal on Scientific Computing* 30.1 (2008), pp. 232–245.
- [19] Y. Saad. *Iterative methods for sparse linear systems*. SIAM, 2003.
- [20] T.T. Chisholm and D.W. Zingg. “A Jacobian-free Newton–Krylov algorithm for compressible turbulent fluid flows”. In: *Journal of Computational Physics* 228.9 (2009), pp. 3490–3507.