

Subspace Acceleration for a Sequence of Linear Systems and Application to Plasma Simulation

Margherita Guido^{*†} Daniel Kressner^{*} Paolo Ricci[†]

March 29, 2024

Abstract

We present an acceleration method for sequences of large-scale linear systems, such as the ones arising from the numerical solution of time-dependent partial differential equations coupled with algebraic constraints. We discuss different approaches to leverage the subspace containing the history of solutions computed at previous time steps in order to generate a good initial guess for the iterative solver. In particular, we propose a novel combination of reduced-order projection with randomized linear algebra techniques, which drastically reduces the number of iterations needed for convergence. We analyze the accuracy of the initial guess produced by the reduced-order projection when the coefficients of the linear system depend analytically on time. Extending extrapolation results by Demanet and Townsend to a vector-valued setting, we show that the accuracy improves rapidly as the size of the history increases, a theoretical result confirmed by our numerical observations. In particular, we apply the developed method to the simulation of plasma turbulence in the boundary of a fusion device, showing that the time needed for solving the linear systems is significantly reduced.

1 Introduction

The numerical solution of time-dependent partial differential equations (PDEs) often leads to sequences of linear systems of the form

$$A(t_i)\mathbf{x}(t_i) = \mathbf{b}(t_i) \quad i = 0, 1, 2, \dots, \quad (1)$$

where $t_0 < t_1 < t_2 < \dots$ is a discretization of time t , and both the system matrix $A(t_i) \in \mathbb{R}^{n \times n}$ and the right-hand side $\mathbf{b}(t_i) \in \mathbb{R}^n$ depend on time. Typically, the systems (1) are available only consecutively. Such sequences of

^{*}Ecole Polytechnique Fédérale de Lausanne (EPFL), Institute of Mathematics, 1015 Lausanne, Switzerland (margherita.guido@epfl.ch, daniel.kressner@epfl.ch)

[†]Ecole Polytechnique Fédérale de Lausanne (EPFL), Swiss Plasma Center (SPC), 1015 Lausanne, Switzerland (paolo.ricci@epfl.ch)

linear systems can arise in a number of applications, including implicit time stepping schemes for the solution of PDEs or iterative solutions of non-linear equations and optimization problems. A relevant example is given by time-dependent PDEs solved in presence of algebraic constraints. In this case, even when an explicit time stepping method is used to evolve the nonlinear PDE, the discretization of the algebraic constraints leads to linear systems that need to be solved at every (sub-)timestep. This is the case of the simulation of turbulent plasma dynamics [10], where a linear constraint (Maxwell equations) is imposed upon the plasma dynamics described by a set of non linear fluid or kinetics equations. The linear systems resulting from the discretized algebraic constraints may feature millions of degrees of freedom, hence their solution is often computationally very expensive.

One usually expects that the linear system (1) changes slowly in subsequent time steps. This work is focused on exploiting this property to accelerate iterative solvers, such as CG [17] for symmetric positive definite matrices and GMRES [24] for general matrices. An obvious way to do so is to supply the iterative solver for the timestep t_{i+1} with the solution of (1) at timestep t_i , as initial guess. As a more advanced technique, in the context of Krylov subspace methods, *subspace recycling methods* [26] such as GCROT [7] and GMRES-DR [21] have been proposed. Such methods have been developed in the case of a single linear system, to enrich the information when restarting the iterative solver. The idea behind is often to accelerate the convergence by suppressing parts of the spectrum of the matrix, including the corresponding approximate invariant subspace in the Krylov minimization subspace. GCROT and GMRES-DR have then been adapted to sequences of linear systems in [22], recycling selected subspaces from one system to the next. For this class of methods to be efficient, it is necessary that the sequence of matrices undergoes local changes only, that is, the difference $A(t_{i+1}) - A(t_i)$ is computationally cheap to apply. For example, one can expect this difference matrix to be sparse when time dependence is restricted to a small part of the computational domain, e.g., through time-dependent boundary conditions. We refer to [26] for a more complete survey of subspace recycling methods and their applications. In [5], subspace recycling was combined with goal-oriented POD (Proper Orthogonal Decomposition) in order to limit the size of the subspaces involved in an augmented CG approach. Simplifications occur when the matrices $A(t_i)$ are actually a fixed matrix A shifted by different scalar multiples of the identity matrix, because Krylov subspaces are invariant under such shifts. In the context of subspace recycling, this property has been exploited in, e.g., [27], and in [25] it is shown how a smoothly varying right-hand side can be incorporated.

When $A(t_i)$ and $\mathbf{b}(t_i)$ in (1) are samples of smooth matrix/vector-valued functions, one expects that the subspace of the previously computed solutions contains a very good approximation of the current one. This can be exploited to construct a better initial guess, either explicitly through (polynomial) extrapolation, or implicitly through projection techniques. Examples of the extrapolation approach include polynomial POD extrapolation [14], weighted group extrapolation methods [30] and a stabilized, least-squares polynomial extrapolation

method [1], for the case that only the right-hand side evolves in time. For the same setting, projection techniques have been introduced by Fischer [11]. Following this first work, several approaches have been developed to extract an initial guess from the solution of a reduced-order model, constructed from projecting the problem to a low-dimensional subspace spanned by previous solutions. In [28], such an approach is applied to fully implicit discretizations of nonlinear evolution problems, while [20] applies the same idea to the so called IMPES scheme used for simulating two-phase flows through heterogeneous porous media.

In this paper, we develop a new projection technique for solving sequences of linear systems that combines projection with randomized linear algebra techniques, leading to considerably reduced cost. Moreover, a novel convergence analysis of the algorithm is carried out to show its efficiency. This is also proved numerically by applying the algorithm to the numerical simulation of turbulent plasma in the boundary of a fusion device.

The rest of this paper is organized as follows. In Section 2, we first discuss general subspace acceleration techniques based on solving a projected linear system and then explain how randomized techniques can be used to speed up existing approaches. In Section 3, a convergence analysis of these subspace acceleration techniques is presented. In Section 4 we first discuss numerical results for a test case to demonstrate the improvements that can be attained by the new algorithm in a somewhat idealistic setting. In Section 5 our algorithm is applied to large-scale turbulent simulation of plasma in a tokamak, showing a significant reduction of computational time.

2 Algorithm

The algorithm proposed in this work for accelerating the solution of the sequence of linear systems (1) uses randomized techniques to lower the cost of a POD-based strategy, such as the one proposed in [20]. Recall that we aim at solving the linear systems $A(t_i)\mathbf{x}(t_i) = \mathbf{b}(t_i)$ consecutively for $i = 0, 1, \dots$. We make no assumption on the symmetry of $A(t_i) \in \mathbb{R}^{n \times n}$ and thus GMRES is an appropriate choice for solving each linear system. Supposing that, at the i th timestep, M previous solutions are available, we arrange them into the history matrix

$$X = [\mathbf{x}(t_{i-M}) \mid \dots \mid \mathbf{x}(t_{i-1})] \in \mathbb{R}^{n \times M}.$$

where the notation on the right-hand side indicates the concatenation of columns. Instead of using the complete history, which may contain redundant information, one usually selects a subspace $\mathcal{S} \subset \text{span}(X)$ of lower dimension $m \leq M$. Then, the initial guess for the i th linear system is obtained from choosing the element of $\mathcal{S}$ that minimizes the residual:

$$\min_{\mathbf{s} \in \mathcal{S}} \|A(t_i)\mathbf{s} - \mathbf{b}(t_i)\|_2 = \min_{\mathbf{z} \in \mathbb{R}^m} \|A(t_i)Q\mathbf{z} - \mathbf{b}(t_i)\|_2,$$

where the columns of $Q \in \mathbb{R}^{n \times m}$ contain an orthonormal basis of $\mathcal{S}$. We use $\|\cdot\|_2$ to denote the Euclidean norm for vectors and the spectral norm for matrices. The described approach is summarized in Algorithm 1, which is a template that needs to be completed by an appropriate choice of the subspace $\mathcal{S}$, in Sections 2.1 and 2.2.

Algorithm 1 Solution of i th linear system $A(t_i)\mathbf{x}(t_i) = \mathbf{b}(t_i)$

Require: History of M solutions $\{\mathbf{x}(t_{i-M}), \dots, \mathbf{x}(t_{i-1})\}$

- 1: $X = [\mathbf{x}(t_{i-M}) \mid \dots \mid \mathbf{x}(t_{i-1})]$
 - 2: Generate $Q \leftarrow$ orthonormal basis for $\mathcal{S} \subseteq \text{span}(X)$, $\dim(\mathcal{S}) = m \leq M$
 - 3: Compute $\mathbf{s}^* = \underset{\mathbf{z} \in \mathbb{R}^m}{\text{argmin}} \|A(t_i)Q\mathbf{z} - \mathbf{b}(t_i)\|_2 \in \mathcal{S}$
 - 4: Solve $A(t_i)\mathbf{x}(t_i) = \mathbf{b}(t_i)$ using GMRES with initial guess $\mathbf{s}^* \in \mathcal{S}$
-

If the complete history is used, $\mathcal{S} = \text{span}(X)$, then computing Q via a QR decomposition [13], as required in Step 2, costs $\mathcal{O}(M^2n)$ operations. In addition, setting up the linear least-squares problem in Step 3 of Algorithm 1 requires M (sparse) matrix-vector products in order to compute $A(t_i)Q$. The standard approach for solving the linear least-squares problem proceeds through the QR decomposition of that matrix and costs another $\mathcal{O}(M^3 + M^2n)$ operations. This strong dependence of the cost on M effectively forces a rather small choice of M , neglecting relevant components of the solutions that could be contained in older solutions only. In the following, we discuss two strategies to overcome this problem.

2.1 Proper Orthogonal Decomposition

An existing strategy [20] to arrive at a low-dimensional subspace $\mathcal{S} \subset \text{span}(X)$ uses a POD approach [19] and computes the orthonormal basis Q for $\mathcal{S}$ through a truncated SVD (Singular Value Decomposition) of X ; see Algorithm 2. Note that only the first m left singular vectors $\Psi_1, \dots, \Psi_m$ need to be computed in Step 2.

Algorithm 2 Method 1 (POD) to generate basis $Q = Q_{\text{POD}}$

Require: History of M solutions $\{\mathbf{x}(t_{i-M}), \dots, \mathbf{x}(t_{i-1})\}$

- 1: $X = [\mathbf{x}(t_{i-M}) \mid \dots \mid \mathbf{x}(t_{i-1})]$
 - 2: Compute SVD of X : $[\Psi, \Sigma, \Phi] = \text{svd}(X)$
 - 3: $Q_{\text{POD}} = [\Psi_1 \mid \dots \mid \Psi_m] \in \mathbb{R}^{n \times m}$
-

Thanks to basic properties of the SVD, the basis Q_{POD} enjoys the following

optimality property [29]:

$$\|(I - Q_{\text{POD}}Q_{\text{POD}}^T)X\|_F^2 = \sum_{k=m+1}^M \sigma_k^2 = \min_{\substack{Q \in \mathbb{R}^{n \times m} \\ Q^T Q = I_m}} \|(I - QQ^T)X\|_F^2, \quad (2)$$

where $\|\cdot\|_F$ denotes the Frobenius norm and $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_M \geq 0$ are the singular values of X . In words, the choice Q_{POD} minimizes the error of orthogonally projecting the columns of X onto an m -dimensional subspace. The relation to the singular values of X established in (2) also allows one to choose m adaptively, e.g., by choosing m such that most of the variability in the history matrix X is captured.

At every time step, the history matrix X gets modified by removing its first column and appending a new last column. The most straightforward implementation of Algorithm 2 would compute the SVD needed of Step 2 from scratch at every time step, leading to a complexity of $\mathcal{O}(nM^2)$ operations. In principle, SVD updating techniques, such as the ones presented in [4] and [6], could be used to reduce this complexity to $\mathcal{O}(mn + m^3)$ for every time step. However, in the context of our application, there is no need to update a complete SVD (in particular, the right singular vectors are not needed) and the randomized techniques discussed in the next section seem to be preferable.

2.2 Randomized range finder

In this section, an alternative to the POD method (Algorithm 1) for generating the low-dimensional subspace $\mathcal{S} \subset \text{span}(X)$ is presented, relying on randomized techniques. The randomized SVD from [16] applied to the $n \times M$ history matrix X proceeds as follows. First, we draw an $M \times m$ Gaussian random matrix Z , that is, the entries of Z are independent and identically distributed (i.i.d) standard normal variables. Then the so-called *sketch*

$$\Omega = XZ = [\mathbf{x}(t_{i-M}) \mid \dots \mid \mathbf{x}(t_{i-1})] Z \quad (3)$$

is computed, followed by a reduced QR decomposition $\Omega = QR$. This only involves the $n \times m$ matrix Ω , which for $m \ll M$ is a significant advantage compared to Algorithm 2, which requires the SVD of an $n \times M$ matrix. The described procedure is contained in lines 2–4 and 11 of Algorithm 3 below.

According to [16, Theorem 10.5], the expected value (with respect to Z) of the error returned by the randomized SVD satisfies

$$\mathbb{E}\|(I - QQ^T)X\|_F \leq \left(1 + \frac{r}{p-1}\right)^{1/2} \left(\sum_{k>r} \sigma_k^2\right)^{1/2}, \quad (4)$$

where we partition $m = r + p$ for a small oversampling parameter $p \geq 2$. Also, the tail bound from [16, Theorem 10.7] implies that it is highly unlikely that the error is much larger than the upper bound (4). Comparing (4) with the error (2), we see that the randomized method is only a factor $\sqrt{2}$ worse than

the optimal basis of roughly *half* the size produced by POD. As we also see in our experiments of Section 4, this bound is quite pessimistic and usually the randomized SVD performs nearly as good as POD using bases of the *same* size.

Algorithm 3 Method 2 (Randomized Range Finder) to generate basis Q

Require: History of M solutions $\{\mathbf{x}(t_{i-M}), \dots, \mathbf{x}(t_{i-1})\}$.

Optional: $\mathbf{x}(t_{i-M-1})$, matrices Ω and Z from previous time step (see (5))

- 1: **if** Ω is computed from scratch **then**
 - 2: $X = [\mathbf{x}(t_{i-M}) \mid \dots \mid \mathbf{x}(t_{i-1})] \in \mathbb{R}^{n \times M}$
 - 3: Draw Gaussian random matrix $Z = [\mathbf{z}_1 \mid \dots \mid \mathbf{z}_M]^T \in \mathbb{R}^{M \times m}$
 - 4: $\Omega = XZ \in \mathbb{R}^{n \times m}$
 - 5: **else**
 - 6: $\Omega = \Omega - \mathbf{x}(t_{i-M-1})\mathbf{z}_1$ $\triangleright \Omega$ is updated
 - 7: $\mathbf{z}_k = \mathbf{z}_{k+1} \quad k = 1, \dots, M-1$
 - 8: Draw new Gaussian random vector $\mathbf{z}_M \in \mathbb{R}^m$
 - 9: $\Omega = \Omega + \mathbf{x}(t_{i-1})\mathbf{z}_M^T$ $\triangleright \Omega$ is updated
 - 10: **end if**
 - 11: $[Q, R] = \text{reduced QR of } \Omega$
-

Instead of performing the randomized SVD from scratch in every timestep, one can easily exploit the fact that only a small part of the history matrix is modified. To see this, let us consider the sketch from the previous timestep:

$$\Omega_{\text{prev}} = [\mathbf{x}(t_{i-M-1}) \mid \dots \mid \mathbf{x}(t_{i-2})] Z^{\text{prev}}. \quad (5)$$

Comparing with (3), we see that the sketch Ω of the current timestep is obtained by removing the contribution from the solution $\mathbf{x}(t_{i-M-1})$ and adding the contribution of $\mathbf{x}(t_{i-1})$. The removal is accomplished in line 6 of Algorithm 3 by a rank-one update:

$$\Omega_{\text{prev}} - \mathbf{x}(t_{i-M-1})\mathbf{z}_1^{\text{prev}} = [\mathbf{0} \mid \mathbf{x}(t_{i-M}) \mid \dots \mid \mathbf{x}(t_{i-2})] Z^{\text{prev}}.$$

By a cyclic permutation, we can move the zero column to the last column, $[\mathbf{x}(t_{i-M}) \mid \dots \mid \mathbf{x}(t_{i-2}) \mid \mathbf{0}]$, updating Z as in line 7 of Algorithm 3. Finally, the contribution of the latest solution is incorporated by adding the rank-one matrix $\mathbf{x}(t_{i-1})\mathbf{z}_M^T$, where $\mathbf{z}_M \in \mathbb{R}^m$ is a newly generated Gaussian random vector that is stored in the last row of Z . Under the (idealistic) assumption that all solutions are exactly computed (and hence deterministic), the described progressive updating procedure is mathematically equivalent to computing the randomized SVD from scratch. In particular, the error bound (4) continues to hold.

Lines 6–9 of Algorithm 3 require $\mathcal{O}(nm)$ operations. When using standard updating procedures for QR decomposition [13], line 11 has the same complexity.

This compares favorably with the $\mathcal{O}(nM^2)$ operations needed by Algorithm 2 per timestep.

When performing the progressive update of Ω over many timesteps, one can encounter numerical issues due to numerical cancellation in the repeated subtraction and addition of contributions to the sketch matrix. To avoid this, the progressive update is carried out only for a fixed number of timesteps, after which a new random matrix Z is periodically generated and Ω is computed from scratch.

3 Convergence Analysis

We start our convergence analysis of the algorithms from the preceding section by considering analytical properties of the history matrix $X = [\mathbf{x}(t_{i-M}) \mid \cdots \mid \mathbf{x}(t_{i-1})]$. After reparametrization, we may assume without loss of generality that each of the past timesteps is contained in the interval $[-1, 1]$:

$$-1 = t_{i-M} < \cdots < t_{i-1} = 1.$$

For notational convenience, we define

$$X \equiv X(\mathbf{t}) := [\mathbf{x}(t_{i-M}) \mid \cdots \mid \mathbf{x}(t_{i-1})], \quad \mathbf{t} = [t_{i-M}, \cdots, t_{i-1}], \quad (6)$$

where $\mathbf{x}(t)$ satisfies the (parametrized) linear system

$$A(t)\mathbf{x}(t) = \mathbf{b}(t), \quad A : [-1, 1] \rightarrow \mathbb{R}^{n \times n}, \quad \mathbf{b} : [-1, 1] \rightarrow \mathbb{R}^n, \quad (7)$$

that is, each entry of A and $\mathbf{b}$ is a scalar function on the interval $[-1, 1]$. Indeed, for the convergence analysis, we assume that each linear system of the sequence in (1) is obtained by sampling the parametrized system in (7) in $t_i \in [-1, 1]$. In many practical applications, like the one described in Section 5, the time dependence in (7) arises from time-dependent coefficients in the underlying PDEs. Frequently, this dependence is real analytic, which prompts us to make the following smoothness assumption on A , $\mathbf{b}$.

Assumption 1. *Consider the open Bernstein ellipse $E_\rho \subset \mathbb{C}$ for $\rho > 1$, that is, the open ellipse with foci ± 1 and semi-minor/-major axes summing up to ρ . We assume that $A : [-1, 1] \rightarrow \mathbb{C}^{n \times n}$ and $\mathbf{b} : [-1, 1] \rightarrow \mathbb{C}^n$ admit extensions that are analytic on E_ρ and continuous on $\bar{E}_\rho$ (the closed Bernstein ellipse), such that $A(t)$ is invertible for all $t \in \bar{E}_\rho$. In particular, this implies that $\mathbf{x}(t) = A^{-1}(t)\mathbf{b}(t)$ is analytic on E_ρ and $\kappa_\rho := \max_{t \in \partial E_\rho} \|\mathbf{x}(t)\|_2$ is finite.*

3.1 Compressibility of the solution time history

The effectiveness of POD-based algorithms relies on the compressibility of the solution history, that is, the columns of X can be well approximated by an m -dimensional subspace with $m \ll M$. According to (2), this is equivalent to stating that the singular values of X decrease rapidly to zero. Indeed, this

property is implied by Assumption 1 as shown by the following result, which was stated in [18] in the context of low-rank methods for solving parametrized linear systems.

Theorem 2 ([18, Theorem 2.4]). *Under Assumption 1, the k th largest singular value σ_k of the history matrix $X(\mathbf{t})$ from (6) satisfies*

$$\sigma_k \leq \frac{2\rho\kappa_\rho\sqrt{M}}{1-\rho^{-1}}\rho^{-k}.$$

Combined with (2), Theorem 2 implies that the POD basis $Q_{\text{POD}} \in \mathbb{R}^{n \times m}$ satisfies the error bound

$$\|(I - Q_{\text{POD}}Q_{\text{POD}}^T)X\|_F^2 \leq \frac{4\rho^2\kappa_\rho^2M}{(1-\rho^{-1})^2}(\rho^{-(m+1)} - \rho^{-(M+1)}).$$

3.2 Quality of prediction without compression

Algorithm 1 determines the initial guess $\mathbf{s}^*$ for the next time step $t_i > t_{i-1} = 1$ by solving the minimization problem

$$\mathbf{s}^* = \underset{\mathbf{s} \in \mathcal{S}}{\operatorname{argmin}} \|A(t_i)\mathbf{s} - \mathbf{b}(t_i)\|_2. \quad (8)$$

In this section, we will assume, additionally to Assumption (1), that $\mathcal{S} = \operatorname{span}(X(\mathbf{t}))$, that is, $X(\mathbf{t})$ is not compressed. Our analysis focuses on uniform timesteps $\mathbf{t}_{\text{equi}} = [t_{i-M}, \dots, t_{i-1}]$ defined by

$$t_{i-M} = -1, t_{i-M+1} = -1 + \Delta t, \dots, t_{i-2} = 1 - \Delta t, t_{i-1} = 1, \quad \Delta t = 2/(M-1).$$

Note that the next timestep $t_i = 1 + \Delta t$ satisfies $t_i \in E_\rho$ if and only if $\rho > t_i + \sqrt{t_i^2 - 1} \approx 1 + \sqrt{2}\Delta t$. The following result shows how the quality of the initial guess rapidly improves (at a square root exponential rate, compared to the exponential rate of Theorem 2) as M , the number of previous time steps in the history, increases.

Theorem 3. *Under Assumption (1), the initial guess constructed by Algorithm 1 with $\mathcal{S} = \operatorname{span}(X)$ satisfies the error bound*

$$\|A(t_i)\mathbf{s}^* - \mathbf{b}(t_i)\|_2 \leq 2\|A(t_i)\|_2\kappa_\rho \left[\frac{1}{1-r} + \frac{C(M, R)\rho}{(\rho-1)\sqrt{\rho^2r^2-1}} \right] r^{R+1},$$

with $C(M, R) = 5\sqrt{5}\sqrt{2R+1}\sqrt{M}/\sqrt{2(M-1)}$, for any $R \leq \frac{1}{2}\sqrt{M-1}$, $r = (t_i + \sqrt{t_i^2 - 1})/\rho < 1$.

3.2.1 Proof of Theorem 3

The rest of this section is concerned with the proof of Theorem 3. We establish the result by making a connection to vector-valued polynomial extrapolation and

extending results by Demanet and Townsend [8] on polynomial extrapolation to the vector-valued setting.

Let $\mathbb{P}_R \subset \mathbb{R}^n[t]$ denote the subspace of vector-valued polynomials of length n and degree at most R for some $R \leq M - 1$. We recall that any $\mathbf{v} \in \mathbb{P}_R$ takes the form $\mathbf{v}(t) = \mathbf{v}_0 + \mathbf{v}_1 t + \cdots + \mathbf{v}_R t^R$ for constant vectors $\mathbf{v}_0, \dots, \mathbf{v}_R \in \mathbb{R}^n$. Equivalently, each entry of $\mathbf{v}$ is a (scalar) polynomial of degree at most R . In our analysis we consider vector-valued polynomials of the particular form

$$\mathbf{p}(t) = X(\mathbf{t}_{\text{equi}})\mathbf{y}(t), \quad (9)$$

for a vector-valued polynomial $\mathbf{y}(t)$ of length M . A key observation is that the evaluation of $\mathbf{p}$ in the next timestep t_i satisfies $\mathbf{p}(t_i) \in \text{span}(X(\mathbf{t}_{\text{equi}})) = \mathcal{S}$. According to (8), $\mathbf{s}^*$ minimizes the residual over $\mathcal{S}$. Hence, the residual can only increase when we replace $\mathbf{s}^*$ by $\mathbf{p}(t_i)$ in

$$\begin{aligned} \|A(t_i)\mathbf{s}^* - \mathbf{b}(t_i)\|_2 &\leq \|A(t_i)\mathbf{p}(t_i) - \mathbf{b}(t_i)\|_2 \\ &\leq \|A(t_i)\|_2 \|\mathbf{p}(t_i) - \mathbf{x}(t_i)\|_2. \end{aligned} \quad (10)$$

Thus, it remains to find a polynomial of the form (9) for which we can establish convergence of the extrapolation error $\|\mathbf{p}(t_i) - \mathbf{x}(t_i)\|_2$. For this purpose, we will choose $\mathbf{p}_R \in \mathbb{P}_R$ to be the least-squares approximation of the M function samples contained in $X(\mathbf{t}_{\text{equi}})$:

$$\mathbf{p}_R := \underset{\mathbf{p} \in \mathbb{P}_R}{\text{argmin}} \|X(\mathbf{t}_{\text{equi}}) - P(\mathbf{t}_{\text{equi}})\|_F, \quad P(\mathbf{t}_{\text{equi}}) = [\mathbf{p}(t_{i-M}) \mid \cdots \mid \mathbf{p}(t_{i-1})]. \quad (11)$$

We will represent the entries of $\mathbf{p}_R$ in the Chebyshev polynomial basis:

$$\mathbf{p}_R(t) = q_0(t)\mathbf{c}_{0,p} + q_1(t)\mathbf{c}_{1,p} + \cdots + q_R(t)\mathbf{c}_{R,p}, \quad (12)$$

where $\mathbf{c}_{k,p} \in \mathbb{R}^n$ and q_k denotes the Chebyshev polynomial of degree k , that is, $q_k(t) = \cos(k \cos^{-1} t)$ for $t \in [-1, 1]$. Setting

$$C_p = [\mathbf{c}_{0,p} \mid \cdots \mid \mathbf{c}_{R,p}] \in \mathbb{R}^{n \times (R+1)}, \quad \mathbf{q}_R(t) = [q_0(t), \dots, q_R(t)]^T, \quad (13)$$

we can express (12) more compactly as $\mathbf{p}_R(t) = C_p \mathbf{q}_R(t)$. Thus,

$$P_R(\mathbf{t}_{\text{equi}}) = C_p Q_R(\mathbf{t}_{\text{equi}}), \quad Q_R(\mathbf{t}_{\text{equi}}) = [\mathbf{q}_R(t_1) \mid \cdots \mid \mathbf{q}_R(t_M)].$$

In view of (11), the matrix of coefficients C_p is determined by minimizing $\|X(\mathbf{t}_{\text{equi}}) - C_p Q_R(\mathbf{t}_{\text{equi}})\|_F$. Because $R \leq M - 1$, the matrix $Q_R(\mathbf{t}_{\text{equi}})$ has full row rank and thus the solution of this least-squares problem is given by $C_p = X(\mathbf{t}_{\text{equi}})Q_R(\mathbf{t}_{\text{equi}})^\dagger$ with $Q_R(\mathbf{t}_{\text{equi}})^\dagger = Q_R(\mathbf{t}_{\text{equi}})^T (Q_R(\mathbf{t}_{\text{equi}})Q_R(\mathbf{t}_{\text{equi}})^T)^{-1}$. In summary, we obtain that

$$\mathbf{p}_R(t) = C_p \mathbf{q}_R(t) = X(\mathbf{t}_{\text{equi}})Q_R(\mathbf{t}_{\text{equi}})^\dagger \mathbf{q}_R(t), \quad (14)$$

which is of the form (9) and thus contained in $\text{span}(X(\mathbf{t}_{\text{equi}}))$, as desired.

In order to analyze the convergence of $\mathbf{p}_R(t)$, we relate it to Chebyshev polynomial interpolation of $\mathbf{x}$. The following lemma follows from classical approximation theory, see, e.g., [18, Lemma 2.2].

Lemma 4. Let $\mathbf{q}_R(t) \in \mathbb{R}^{R+1}$ be defined as in (13), containing the Chebyshev polynomials up to degree R . Under Assumption 1 there exists an approximation of the form

$$\mathbf{x}_R(t) = C_x \mathbf{q}_R(t), \quad C_x = [\mathbf{c}_{0,x}, \mathbf{c}_{1,x}, \dots, \mathbf{c}_{R,x}] \in \mathbb{R}^{n \times (R+1)},$$

such that $\|\mathbf{c}_{k,x}\|_2 \leq 2\kappa_\rho \rho^{-k}$ and

$$\max_{t \in [-1,1]} \|\mathbf{x}_R(t) - \mathbf{x}(t)\|_2 \leq \frac{2\kappa_\rho}{\rho - 1} \rho^{-R}.$$

Following the arguments in [8] for scalar functions, Lemma 4 allows us to estimate the extrapolation error for $\mathbf{p}_R(t)$ if $R \sim \sqrt{M}$.

Theorem 5. Suppose that Assumption 1 holds and $R \leq \frac{1}{2}\sqrt{M-1}$. Then the vector-valued polynomial $\mathbf{p}_R \in \mathbb{P}_R$ defined in (14) satisfies for every $t \in (1, (\rho + \rho^{-1})/2)$ the error bound

$$\|\mathbf{x}(t) - \mathbf{p}_R(t)\|_2 \leq 2\kappa_\rho \left[\frac{1}{1-r} + \frac{C(M,R)\rho}{(\rho-1)\sqrt{\rho^2 r^2 - 1}} \right] r^{R+1},$$

with $r = (t + \sqrt{t^2 - 1})/\rho < 1$ and $C(M,R)$ defined as in Theorem 3.

Proof. Letting $\mathbf{x}_R$ be the polynomial from Lemma 4, we write

$$\begin{aligned} \|\mathbf{x}(t) - \mathbf{p}_R(t)\|_2 &\leq \|\mathbf{x}(t) - \mathbf{x}_R(t)\|_2 + \|\mathbf{x}_R(t) - \mathbf{p}_R(t)\|_2 \\ &= \left\| \sum_{k=R+1}^{\infty} \mathbf{c}_{k,x} q_k(t) \right\|_2 + \|(C_x - C_p) \mathbf{q}_R(t)\|_2 \\ &\leq \sum_{k=R+1}^{\infty} \|\mathbf{c}_{k,x}\|_2 |q_k(t)| + \|C_x - C_p\|_2 \|\mathbf{q}_R(t)\|_2. \end{aligned} \quad (15)$$

To treat the second term in (15), first note that, by definition, we have

$$X_R(\mathbf{t}_{\text{equi}}) = [\mathbf{x}_R(t_{i-M}) \mid \dots \mid \mathbf{x}_R(t_{i-1})] = C_x Q_R(\mathbf{t}_{\text{equi}})$$

and hence $C_x = X_R(\mathbf{t}_{\text{equi}}) Q_R(\mathbf{t}_{\text{equi}})^\dagger$. Setting $\sigma := \sigma_{\min}(Q_R(\mathbf{t}_{\text{equi}})) = 1/\|Q_R(\mathbf{t}_{\text{equi}})^\dagger\|_2$, we obtain

$$\begin{aligned} \|C_x - C_p\|_2 &= \|(X_R(\mathbf{t}_{\text{equi}}) - X(\mathbf{t}_{\text{equi}})) Q_R(\mathbf{t}_{\text{equi}})^\dagger\|_2 \leq \|X_R(\mathbf{t}_{\text{equi}}) - X(\mathbf{t}_{\text{equi}})\|_2 / \sigma \\ &\leq \frac{\sqrt{M}}{\sigma} \cdot \max_{k=1, \dots, M} \|\mathbf{x}_R(t_k) - \mathbf{x}(t_k)\|_2 \leq \frac{\sqrt{M}}{\sigma} \frac{2\kappa_\rho}{\rho - 1} \rho^{-R}, \end{aligned}$$

where we used Lemma 4 in the last inequality. Applying, once more, Lemma 4 to the first term in (15) gives

$$\|\mathbf{x}(t) - \mathbf{p}_R(t)\|_2 \leq 2\kappa_\rho \left[\sum_{k=R+1}^{\infty} \rho^{-k} |q_k(t)| + \frac{\sqrt{M}}{\sigma} \frac{\rho^{-R}}{\rho - 1} \|\mathbf{q}_R(t)\|_2 \right] \quad (16)$$

Because $|q_k(t)| \leq (t + \sqrt{t^2 - 1})^k \leq \rho^k r^k$ for $t > 1$, we have that

$$\|\mathbf{q}_R(t)\|_2^2 \leq \sum_{k=0}^R (\rho r)^{2k} = (\rho r)^{2R} \sum_{k=0}^R (\rho r)^{-2k} \leq \frac{(\rho r)^{2R+2}}{\rho^2 r^2 - 1}. \quad (17)$$

Inserted into (16), this gives

$$\begin{aligned} \|\mathbf{x}(t) - \mathbf{p}_R(t)\|_2 &\leq 2\kappa_\rho \left[\sum_{k=R+1}^{\infty} \rho^{-k} \rho^k r^k + \frac{\sqrt{M}}{\sigma} \frac{\rho^{-R} (\rho r)^{R+1}}{(\rho - 1) \sqrt{\rho^2 r^2 - 1}} \right] \\ &\leq 2\kappa_\rho \left[\frac{1}{1-r} + \frac{\sqrt{M}\rho}{\sigma(\rho - 1) \sqrt{\rho^2 r^2 - 1}} \right] r^{R+1}. \end{aligned}$$

The proof is completed by inserting the lower bound

$$\sigma = \sigma_{\min}(Q_R(\mathbf{t}_{\text{equi}})) \geq \frac{\sqrt{2}}{5\sqrt{5}} \frac{\sqrt{M-1}}{\sqrt{2R+1}}, \quad (18)$$

which holds when $R \leq \frac{1}{2}\sqrt{M-1}$ according to [8, Theorem 4]. $\square$

Using Theorem 5 with $t = t_i$ and inserting the result in (10), we have proven the statement of Theorem 3.

3.3 Optimality of the prediction with compression

When the matrix $X(\mathbf{t})$ is compressed via POD (Algorithm 2) or the randomized range finder (Algorithm 3), the orthonormal basis $Q \in \mathbb{R}^{n \times m}$ used in Algorithm 1 spans a lower-dimensional subspace $\mathcal{S} \subseteq \text{span}(X)$.

Corollary 6. *Suppose that Algorithm 1 is used with an orthonormal basis satisfying $\|(QQ^T - I)X(\mathbf{t}_{\text{equi}})\|_2 \leq \varepsilon$ for some tolerance $\varepsilon > 0$. Under Assumption 1, the initial guess $\mathbf{s}^*$ constructed by the algorithm satisfies the error bound*

$$\|A(t_i)\mathbf{s}^* - \mathbf{b}(t_i)\|_2 \leq 2\|A(t_i)\|_2 \kappa_\rho \left[\frac{1}{1-r} + \frac{C(M, R)\rho}{\sqrt{\rho^2 r^2 - 1}} \left(\frac{1}{\rho - 1} + \frac{\varepsilon \rho^R}{2\sqrt{M}\kappa_\rho} \right) \right] r^{R+1}$$

for any $R \leq \frac{1}{2}\sqrt{M-1}$.

Proof. Let $\mathbf{p}_R(t) = X(\mathbf{t}_{\text{equi}})Q_R(\mathbf{t}_{\text{equi}})^\dagger \mathbf{q}_R(t)$ be the polynomial constructed in (14). Using that $\mathbf{s}^*$ satisfies the minimization problem (8) and $QQ^T \mathbf{x}(t_i) \in \mathcal{S} = \text{span}(Q)$, we obtain:

$$\begin{aligned} \|A(t_i)\mathbf{s}^* - \mathbf{b}(t_i)\|_2 &\leq \|A(t_i)QQ^T \mathbf{x}(t_i) - \mathbf{b}(t_i)\|_2 \\ &\leq \|A(t_i)\|_2 [\|(QQ^T - I)(\mathbf{x}(t_i) - \mathbf{p}_R(t_i))\|_2 \\ &\quad + \|(QQ^T - I)\mathbf{p}_R(t_i)\|_2] \\ &\leq \|A(t_i)\|_2 [\|\mathbf{x}(t_i) - \mathbf{p}_R(t_i)\|_2 \\ &\quad + \|(QQ^T - I)X(\mathbf{t}_{\text{equi}})Q_R(\mathbf{t}_{\text{equi}})^\dagger \mathbf{q}_R(t_i)\|_2]. \end{aligned}$$

The first term is bounded using Theorem 5 with $t = t_i$. For the second term, we use the bound in (17) on $\|\mathbf{q}_R(t_{M+1})\|$ to obtain

$$\begin{aligned} \|(QQ^T - I)X(\mathbf{t}_{\text{equi}})Q_R(\mathbf{t}_{\text{equi}})^\dagger \mathbf{q}_R(t)\|_2 &\leq \|(QQ^T - I)X(\mathbf{t}_{\text{equi}})\|_2 \|\mathbf{q}_R(t_{M+1})\|_2 / \sigma \\ &\leq \frac{\varepsilon(\rho r)^{R+1}}{\sigma \sqrt{\rho^2 r^2 - 1}}, \end{aligned}$$

with $\sigma := \sigma_{\min}(Q_R(\mathbf{t}_{\text{equi}}))$. The proof is completed using the lower bound (18) on $\mathbf{x}\sigma$. $\square$

4 Numerical results: Test Case

To test the subspace acceleration algorithms proposed in Section 2, we first consider a simplified setting, an elliptic PDE with an explicitly given time- and space-dependent coefficient $a(\mathbf{x}, t)$ and source term $g(\mathbf{x}, t)$:

$$\begin{cases} \nabla \cdot (a(\mathbf{x}, t) \nabla f(\mathbf{x}, t)) = g(\mathbf{x}, t) & \text{in } \Omega \\ f(\mathbf{x}, t) = 0 & \text{on } \partial\Omega \end{cases} \quad (19)$$

We consider the domain $\Omega = [0, 1]^2 \subset \mathbb{R}^2$ and discretize (19) on a uniform two-dimensional Cartesian grid using a centered finite difference scheme of order 4. This leads to a linear system for the vector of unknowns $\mathbf{f}(t)$, for which both the matrix and the right-hand side depend on t :

$$A(t)\mathbf{f}(t) = \mathbf{g}(t). \quad (20)$$

We discretize the time variable on the interval $[t_0, t_f]$ with a uniform timestep Δt on N_t points, such that $t_f = t_0 + N_t \Delta t$. Evaluating (20) in these N_t instants, we obtain a sequence of linear systems of the same type as (1).

We set $a(\mathbf{x}, t) = \exp[-(x-0.5)^2 - (y-0.5)^2] \cos(tx) + 2.1$ and choose the right-hand side $g(\mathbf{x}, t)$ such that

$$f(\mathbf{x}, t) = \sin(4\pi y t) \sin(15\pi x t) \left[1 + \sin(15\pi x t) \cos(3\pi y t) \exp[(x-0.5)^2 + (y-0.5)^2 - 0.25^2] \right]$$

is the exact solution of (19). The tests are performed using MATLAB 2023a on an M1 MacbookPro. We employ GMRES as iterative solver for the linear system, with tolerance 10^{-7} and incomplete LU factorization as preconditioner. We start the simulations at $t_0 = 2.3$ s and perform $N_t = 200$ timesteps.




Figure 1: GMRES iterations per timestep when solving equation (20) with different initial guesses.

The results reported in Figure 1 use a spatial grid of dimension 100×100 , leading to linear systems of size $n = 10000$. Different values of M , the number of previous solutions retained in the history matrix X , and m , the dimension of the reduced-order model, were tested. We found that the choices $M = 20$, $m = 10$ and $M = 35$, $m = 20$ lead to good performance for $\Delta t = 10^{-5}$ and $\Delta t = 10^{-3}$, respectively. The baseline is (preconditioned) GMRES with the previous solution used as initial guess; the resulting number of iterations is indicated with the solid blue line (“*Baseline*”) in Figure 1. This is compared to the number of iterations obtained by applying GMRES when Algorithm 1 is employed to compute the initial guess, in combination with both the POD basis in Algorithm 2 (“*POD*” in the graph) and the Randomized Range Finder in Algorithm 3 (“*RAND*” in the graph). For the Randomized Range Finder algorithm, the matrix Ω is computed from scratch only every 50 timesteps, while in the other timesteps is updated as described in Algorithm 3, resulting in a computationally efficient version of the algorithm. Both the POD and Randomized versions of the acceleration method give a remarkable gain in computational time with respect to the baseline.

When employing $\Delta t = 10^{-5}$, in Figure 1a, the number of iterations computed by the linear solver vanishes most of the time, since the initial residual computed with the new initial guess is already below the tolerance, set to 10^{-7} in this case. It is worth noticing that the new randomized method gives an acceleration comparable to the existing POD one, but it requires a much lower computational cost, as described in Section 2.

The results obtained for larger timesteps, in Figure 1b, are slightly worse, as expected, since it is less easy to predict new solutions using the previous ones when they are further apart in time. Nevertheless, the gain of the acceleration method is still visible, obtaining always less than half iterations with respect to the baseline and adding the solution of a reduced-order system of dimension $m = 20$ only, compared to the full solution of dimension 10000. The resulting

advantage of the new method can indeed be observed in Figure 2, which compares the computational time needed by the solver using the baseline approach with the one obtained by using the new guess (this includes the time employed to compute the guess). The timings showed are the ones needed to produce the results in Figure 1. The time employed by the POD method has not been included since it is significantly higher than the baseline, as predicted by the analysis in Section 2.1.




Figure 2: Computational time per timestep corresponding to Figure 1a and Figure 1b. The average speedup per iteration of the randomized method with respect to the baseline is a factor 9 for $\Delta t = 10^{-5}$ and a factor 10 for $\Delta t = 10^{-3}$.

5 Numerical results: Plasma Simulations

In this Section, we apply the subspace acceleration method to the numerical simulation of plasma turbulence in the outermost plasma region of tokamaks, where the plasma enters in contact with the surrounding external solid walls, resulting in strongly non-linear phenomena occurring on a large range of time and length scales. In this work, we consider GBS (Global Braginskii Solver) [12, 23], a three-dimensional, flux-driven, two-fluid code developed for the simulation of the plasma dynamics in the boundary of a fusion device. GBS implements the Braginskii two-fluid model [3], which describes a quasi-neutral plasma through the conservation of density, momentum, and energy. This results in six coupled three-dimensional time-evolving non-linear equations which evolve the plasma dynamics in Ω , a 3D toroidal domain with rectangular poloidal cross section, as represented in Figure 3. The fluid equations are coupled with Maxwell equations, specifically Poisson and Ampère, elliptic equations for the electromagnetic variables of the plasma. In the limit considered here the elliptic equations reduce to a set of two-dimensional algebraic constraints decoupled along the toroidal direction, therefore to be satisfied independently on each poloidal plane. The




Figure 3: GBS computational domain. The toroidal direction is along φ , the radial direction is along R , and the vertical direction is along Z . The domain consists of N_φ rectangular poloidal planes, each discretized on a $N_R \times N_Z$ Cartesian grid.

differential equations are spatially discretized on a uniform Cartesian grid employing a finite difference method, resulting in a system of differential-algebraic equations of index one [15]:

$$\begin{cases} \partial_t \mathbf{f}(t) = \mathcal{V}(\mathbf{f}(t), \mathbf{x}(t)) & \text{in } \Omega \\ A_k(\mathbf{f}(t)) \mathbf{x}_k(t) = \mathbf{b}_k(\mathbf{f}(t)) & \text{for each } k\text{th poloidal plane} \end{cases} \quad (21)$$

where $\mathcal{V}(\mathbf{f}(t), \mathbf{x}(t))$ is a non-linear, 6-dimensional differential operator and

$$\begin{aligned} \mathbf{x}(t) &= [\mathbf{x}_1(t), \dots, \mathbf{x}_k(t) \dots, \mathbf{x}_{N_Z}(t)] \in \mathbb{R}^{N_R N_\varphi N_Z}, \\ \mathbf{f}(t) &= [\mathbf{f}_1(t), \dots, \mathbf{f}_k(t) \dots, \mathbf{f}_{N_Z}(t)] \in \mathbb{R}^{N_R N_\varphi N_Z} \end{aligned}$$

are the vector of, respectively, the electromagnetic and fluid quantities solved for by GBS, where the solutions of all the N_Z poloidal planes are stacked together. More precisely, the time evolution of the fluid variables, $\mathbf{f}$, is coupled with the set of linear systems $A_k(\mathbf{f}(t)) \mathbf{x}_k(t) = \mathbf{b}_k(\mathbf{f}(t))$ which result from the discretization of Maxwell equations. Indeed, the matrix $A_k \in \mathbb{R}^{N_R N_Z \times N_R N_Z}$ and right-hand side $\mathbf{b}_k \in \mathbb{R}^{N_R N_Z}$ depend on time through $\mathbf{f}$.

In GBS, system (21) is integrated using a Runge-Kutta scheme of order four, on the discrete times $\{t_i\}_{i=1}^{N_t}$, with step-size Δt . Given $\mathbf{f}^i$ and $\mathbf{x}^i$, the value of $\mathbf{f}$ and $\mathbf{x}$ at time t_i , the computation of $\mathbf{f}^{i+1}$, requires performing three intermediate substeps where the quantities $\mathbf{f}^{i+1,j}$ for $j = 1, 2, 3$ are computed. To guarantee the consistency and convergence of the Runge-Kutta integration method [15], the algebraic constraints are solved at every substep, computing $\mathbf{x}_k^{i+1,j}$ for $j = 1, 2, 3$ and for each k -th poloidal plane. As a consequence, the linear systems $A_k(\mathbf{f}(t)) \mathbf{x}_k(t) = \mathbf{b}_k(\mathbf{f}(t))$ are assembled and solved four times for each of the N_φ poloidal planes, to advance the full system (21) by one timestep. Since the timestep Δt is constrained to be small from the stiff nature of the GBS model, the solution of the linear systems is among the most computationally expensive part of GBS simulations.

In GBS, the linear system is solved using GMRES, with the algebraic multi-grid preconditioner *boomerAMG* from the HYPRE library [9], a choice motivated by previous investigations [12]. The subspace acceleration algorithm proposed in Section 2 is implemented in the GBS code and, given the results shown in Section 4, the randomized version of the algorithm is chosen. The results reported are obtained from GBS simulations on one computing node. The poloidal planes of the computational domain are distributed among 16 cores, specifically of type Intel(R) Core i7-10700F CPU at 2.90GHz. GBS is implemented in Fortran 90, and relies on the PETSc library [2] for the linear solver and Intel MPI 19.1 for the parallelization.

We consider the simulation setting described in [12], taking as initial conditions the results of a simulation in a turbulent state. We use a Cartesian grid of size of $N_R = 150$, $N_Z = 300$ and $N_\varphi = 64$, with additional 4 ghost points in the Z and R directions. Therefore, the imposed algebraic constraints result in 64 sequences of linear systems of dimension $N_R N_Z \times N_R N_Z = 46816 \times 46816$. The timestep employed is $\Delta t = 0.7 \times 10^{-5}$. The sequence of linear systems we consider represents the solution of the Poisson equation on one fixed poloidal plane, but the same considerations apply to the discretization of Ampère equation.




Figure 4: Performance of the algorithm applied to the solution of Poisson equation in GBS simulations. The time for the RAND algorithm is on average approximately one fourth of the time for the baseline.

In Figure 4a the number of iterations obtained with the method proposed in Section 2, denoted as “*RAND*” is compared with the ones obtained using the previous step solution as initial guess, depicted in blue as “*Baseline*”. We notice that, employing the acceleration method, the number of GMRES iterations needed for each solution of the linear system is reduced by a factor 2.9, on average, at the cost of computing a solution of an $m \times m$ reduced-order system. In Figure 4b the wall clock time required for the solution of the systems is shown. The baseline approach is compared to the accelerated method, where we also take into account the cost of computing the initial guess. Thanks to the randomized method employed, the process of generating the guess is fast

enough to provide a time speed up of a factor of 6.5 per iteration.

The employed values of $M = 15$, the number of previous solutions retained, and $m = 10$, the dimension of the reduced-order model, are the ones found to give a good balance between the decrease in the number of iterations and the computational cost of the reduced-order model. In Table 1 the results for different values of M and m are reported. It is worth noticing that an average number of GMRES iterations per timestep smaller than one implies that often the initial residual obtained with the initial guess is below the tolerance set for the solver. It is possible to notice that higher values of m lead to very small number of iterations, but the overall time speedup is reduced since the computation of the guess becomes more expensive.

M	m	Average time per timestep [s]	Time speedup	Average GMRES iterations per timestep	Iterations speedup
15	6	0.0452	2.1797	2.0183	2.4743
15	8	0.0347	2.8352	1.1098	4.5
15	10	0.0339	2.9071	0.76	6.552
20	10	0.0358	2.7468	0.7862	6.336
20	15	0.0435	2.2572	0.5031	9.9
30	8	0.0485	2.0255	1.698	2.9328
30	12	0.0375	2.6185	0.3758	13.25
30	15	0.05	1.9633	0.579	8.3371

Table 1: GBS simulations result corresponding to different values of M and m .

The iteration and time speedups are computed on the total of 180 linear systems, with respect to the baseline, that has an average number of 5 GMRES iterations and an average time per timestep of 0.0828 s. The highlighted row corresponds to the best result obtained in terms of time speedup.

6 Conclusions

In this paper, we propose a novel approach for accelerating the solution of a sequence of large-scale linear systems that arises from, e.g., the discretization of time-dependent PDEs. Our method generates an initial guess from the solution of a reduced-order model, obtained by extracting relevant components of previously computed solutions using dimensionality reduction techniques. Starting from an existing POD-like approach, we accelerate the process by employing a randomized algorithm. A convergence analysis is performed, which applies to both approaches, POD and the randomized algorithm and shows how the accuracy of the method increases with the history size. A test case displays how POD leads to a noticeable decrease in the number of iterations, but at the same time a nearly equal decrease is achieved by the cheaper randomized method, that leads to a time speedup per iteration of a factor 9. In real applications such as the plasma simulations described in Section 5, the speedup is more mod-

est, given the stiff nature of the problem which constrains the timestep of the explicit integration method to be very small, but still practically relevant.

Acknowledgements. The authors thank the anonymous reviewers for helpful feedback.

This work has been carried out within the framework of the EUROfusion Consortium, via the Euratom Research and Training Programme (Grant Agreement No 101052200 — EUROfusion) and funded by the Swiss State Secretariat for Education, Research and Innovation (SERI). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union, the European Commission, or SERI. Neither the European Union nor the European Commission nor SERI can be held responsible for them.

Data Availability Statement. The data that support the findings of this study are available upon reasonable request from the authors.

References

- [1] Anthony P. Austin, Noel Chalmers, and Tim Warburton. “Initial guesses for sequences of linear systems in a GPU-accelerated incompressible flow solver”. In: *SIAM J. Sci. Comput.* 43.4 (2021), pp. C259–C289. DOI: 10.1137/20M1368677.
- [2] Satish Balay et al. *PETSc, the portable, extensible toolkit for scientific computation*. Vol. 2. 17. Argonne National Laboratory, 1998.
- [3] S. I. Braginskii. “Transport Processes in a Plasma”. In: *Rev. Plasma Phys.* 1 (Jan. 1965), p. 205.
- [4] Matthew Brand. “Fast low-rank modifications of the thin singular value decomposition”. In: *Linear Algebra Appl.* 415.1 (2006), pp. 20–30. DOI: 10.1016/j.laa.2005.07.021.
- [5] Kevin Carlberg, Virginia Forstall, and Ray Tuminaro. “Krylov-Subspace Recycling via the POD-Augmented Conjugate-Gradient Method”. In: *SIAM Journal on Matrix Analysis and Applications* 37.3 (2016), pp. 1304–1336. DOI: 10.1137/16M1057693.
- [6] Gang Chen, Yangwen Zhang, and Dujin Zuo. *An Incremental SVD Method for Non-Fickian Flows in Porous Media: Addressing Storage and Computational Challenges*. 2023. arXiv: 2308.15409 [math.NA].
- [7] Eric De Sturler. “Truncation strategies for optimal Krylov subspace methods”. In: *SIAM J. Numer. Anal.* 36.3 (1999), pp. 864–889. DOI: 10.1137/S0036142997315950.
- [8] Laurent Demanet and Alex Townsend. “Stable extrapolation of analytic functions”. In: *Found. Comput. Math.* 19.2 (2019), pp. 297–331. DOI: 10.1007/s10208-018-9384-1.

- [9] Robert D Falgout and Ulrike Meier Yang. “hypre: A Library of High Performance Preconditioners”. In: *Comput. Sci. — ICCS 2002*. Ed. by Peter M A Sloot et al. Berlin, Heidelberg: Springer Berlin Heidelberg, 2002, pp. 632–641.
- [10] A. Fasoli et al. “Computational challenges in magnetic-confinement fusion physics”. In: *Nat. Phys.* 12.5 (2016), pp. 411–423. DOI: 10.1038/NPHYS3744.
- [11] Paul F. Fischer. “Projection techniques for iterative solution of $Ax = b$ with successive right-hand sides”. In: *Comput. Methods Appl. Mech. Eng.* 163.1-4 (1998), pp. 193–204. DOI: 10.1016/S0045-7825(98)00012-7.
- [12] M. Giacomini et al. “The GBS code for the self-consistent simulation of plasma turbulence and kinetic neutral dynamics in the tokamak boundary”. In: *J. Comput. Phys.* 463 (2022), p. 111294. DOI: <https://doi.org/10.1016/j.jcp.2022.111294>.
- [13] Gene H. Golub and Charles F. Van Loan. *Matrix computations*. Fourth. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, Baltimore, MD, 2013.
- [14] Leopold Grinberg and George Em Karniadakis. “Extrapolation-based acceleration of iterative solvers: Application to simulation of 3D flows”. In: *Commun. Comput. Phys.* 9.3 (2011), pp. 607–626. DOI: 10.4208/cicp.301109.080410s.
- [15] Ernst Hairer, Christian Lubich, and Michel Roche. *The numerical solution of differential-algebraic systems by Runge-Kutta methods*. Vol. 1409. Lecture Notes in Mathematics. Springer-Verlag, Berlin, 1989, pp. viii+139. DOI: 10.1007/BFb0093947.
- [16] N. Halko, P. G. Martinsson, and J. A. Tropp. “Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions”. In: *SIAM Rev.* 53.2 (2011), pp. 217–288. DOI: 10.1137/090771806.
- [17] M.R. Hestenes and E. Stiefel. “Methods of conjugate gradients for solving linear systems”. In: *J. Res. Natl. Bur. Stand. (1934)*. 49.6 (1952), p. 409. DOI: 10.6028/jres.049.044.
- [18] Daniel Kressner and Christine Tobler. “Low-rank tensor Krylov subspace methods for parametrized linear systems”. In: *SIAM J. Matrix Anal. Appl.* 32.4 (2011), pp. 1288–1316. DOI: 10.1137/100799010.
- [19] K. Kunisch and S. Volkwein. “Control of the Burgers equation by a reduced-order approach using proper orthogonal decomposition”. In: *J. Optim. Theory Appl.* 102.2 (1999), pp. 345–371. DOI: 10.1023/A:1021732508059.
- [20] R. Markovinović and J. D. Jansen. “Accelerating iterative solution methods using reduced-order models as solution predictors”. In: *Int. J. Numer. Methods Eng.* 68.5 (2006), pp. 525–541. DOI: 10.1002/nme.1721.
- [21] Ronald B. Morgan. “Implicitly restarted GMRES and Arnoldi methods for nonsymmetric systems of equations”. In: *SIAM J. Matrix Anal. Appl.* 21.4 (2000), pp. 1112–1135. DOI: 10.1137/S0895479897321362.

- [22] Michael L. Parks et al. “Recycling Krylov subspaces for sequences of linear systems”. In: *SIAM J. Sci. Comput.* 28.5 (2006), pp. 1651–1674. DOI: 10.1137/040607277.
- [23] P. Ricci et al. “Simulation of plasma turbulence in scrape-off layer conditions: The GBS code, simulation results and code validation”. In: *Plasma Phys. Control. Fusion* 54.12 (2012). DOI: 10.1088/0741-3335/54/12/124047.
- [24] Youcef Saad and Martin H. Schultz. “GMRES: A Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear Systems”. In: *SIAM J. Sci. Stat. Comput.* 7.3 (1986), pp. 856–869. DOI: 10.1137/0907058.
- [25] Kirk M. Soodhalter. “Block Krylov Subspace Recycling for Shifted Systems with Unrelated Right-Hand Sides”. In: *SIAM Journal on Scientific Computing* 38.1 (2016), A302–A324. DOI: 10.1137/140998214.
- [26] Kirk M. Soodhalter, Eric de Sturler, and Misha E. Kilmer. “A survey of subspace recycling iterative methods”. In: *GAMM-Mitt.* 43.4 (2020), e202000016, 29. DOI: 10.1002/gamm.202000016.
- [27] Kirk M. Soodhalter, Daniel B. Szyld, and Fei Xue. “Krylov subspace recycling for sequences of shifted linear systems”. In: *Appl. Numer. Math.* 81 (2014), pp. 105–118. DOI: 10.1016/j.apnum.2014.02.006. arXiv: 1301.2650.
- [28] Damien Tromeur-Dervout and Yuri Vassilevski. “Choice of initial guess in iterative solution of series of systems arising in fluid flow simulations”. In: *J. Comput. Phys.* 219.1 (2006), pp. 210–227. DOI: 10.1016/j.jcp.2006.03.014.
- [29] Stefan Volkwein. “Proper orthogonal decomposition: Theory and reduced-order modelling”. In: *Lect. Notes, Univ. Konstanz* 4 (2013), p. 4.
- [30] Shuai Ye et al. “Improving initial guess for the iterative solution of linear equation systems in incompressible flow”. In: *Mathematics* 8.1 (2020), pp. 1–20. DOI: 10.3390/math8010119.