

Relaxed Krylov subspace approximation

Valeria Simoncini*¹ and Daniel B. Szyld²

¹ Dipartimento di Matematica, Università di Bologna, Piazza di Porta S. Donato, 5, I-40127 Bologna, Italy; and also IMATI-CNR, Pavia and CIRSA, Ravenna, Italy

² Department of Mathematics, Temple University (038-16), 1805 N. Broad St., Philadelphia, Pennsylvania 19122-6094, USA

Recent computational and theoretical studies have shown that the matrix-vector product occurring at each step of a Krylov subspace method can be relaxed as the iterations proceed, i.e., it can be computed in a less exact manner, without degradation of the overall performance. In the present paper a general operator treatment of this phenomenon is provided and a new result further explaining its behavior is presented.

Copyright line will be provided by the publisher

1 Introduction

We consider a linear finite n -dimensional operator $v \rightarrow \mathbb{A}(v)$ and a (linear or nonlinear) problem

$$G(\mathbb{A}, x) = 0 \tag{1}$$

to be solved for the vector $x \in \mathbb{C}^n$. Among the problems of interest of the form (1) are those related to the solution of large linear systems, of eigenvalue problems, or to the evaluation of rational functions.

The operator \mathbb{A} is not known exactly or it is too expensive to apply. A typical example is a Schur complement operator of the form $\mathbb{A} = A - C^T D^{-1} E$. Instead, it is possible to employ a relaxed version of \mathbb{A} , i.e., an approximation \mathbb{A}_ϵ such that $\mathbb{A}_\epsilon \rightarrow \mathbb{A}$ for $\epsilon \rightarrow 0$. Here ϵ is a parameter that the user may adjust. The computational cost associated with ϵ increases as $\epsilon \rightarrow 0$. In the example of the Schur complement \mathbb{A}_ϵ can indicate, for example, the approximate solution of a system $Dw = Ez$ within a tolerance of ϵ .

We are interested in efficiently determining the solution x by projecting the problem onto the approximation space

$$\mathcal{K}_m = \mathcal{K}_m(\mathbb{A}_\epsilon, v) = \text{span}\{v, \mathbb{A}_{\epsilon_1}(v), \mathbb{A}_{\epsilon_2}(\mathbb{A}_{\epsilon_1}(v)), \dots, \mathbb{A}_{\epsilon_{m-1}}(\dots \mathbb{A}_{\epsilon_2}(\mathbb{A}_{\epsilon_1}(v)))\}, \quad v \in \mathbb{C}^n,$$

where $\epsilon_k, k = 1, 2, \dots$ are in general different from each other, and we may consider large values of ϵ_k (see further comments below). Note that $\mathcal{K}_m \subseteq \mathcal{K}_{m+1}$, and we assume that \mathcal{K}_m has full dimension m . Clearly, when $\epsilon_k = 0, k = 1, 2, \dots$ (the non-relaxed or exact case) and \mathbb{A} is represented by a matrix A , we have that \mathcal{K} is the usual Krylov subspace $\mathcal{K}_m(A, v) = \text{span}\{v, Av, A^2v, \dots, A^{m-1}v\}$; see, e.g., [2], [4], [11].

Let $V_m = [v_1, \dots, v_m]$ be the matrix of the *orthogonal* basis of $\mathcal{K}_m(\mathbb{A}_\epsilon, v)$ obtained recursively one vector at the time by an Arnoldi procedure, i.e., the orthogonalization of $\mathbb{A}_{\epsilon_k} v_k$ with respect to the previous vectors. In the exact case, we have $v_1 = v/\|v\|$ and

$$\mathbb{A}V_m = V_{m+1}H_{m+1,m}, \tag{2}$$

where $H_{m+1,m}$ is the $(m+1) \times m$ matrix of orthogonalization coefficients, and it is upper Hessenberg. In the general approximated (or relaxed) case, we write the relaxed operator as $\mathbb{A}_{\epsilon_k} = \mathbb{A} + \mathcal{E}_k$, then an orthogonal basis for the relaxed (or inexact) approximation space $\mathcal{K}(\mathbb{A}_\epsilon, v)$ can be obtained by the same process and this gives the relaxed (or inexact) recurrence

$$\mathbb{A}V_m = V_{m+1}H_{m+1,m} - [\mathcal{E}_1 v_1, \mathcal{E}_2 v_2, \dots, \mathcal{E}_m v_m] =: V_{m+1}H_{m+1,m} + F_{m,\epsilon}. \tag{3}$$

Thus, we have that

$$\mathbb{A}_{\epsilon_m} \text{span}\{v, \mathbb{A}_{\epsilon_1}(v), \mathbb{A}_{\epsilon_2}(\mathbb{A}_{\epsilon_1}(v)), \dots, \mathbb{A}_{\epsilon_{m-1}}(\dots \mathbb{A}_{\epsilon_2}(\mathbb{A}_{\epsilon_1}(v)))\} = \text{Range}(V_{m+1}H_{m+1}).$$

Note that neither V_{m+1} nor $H_{m+1,m}$ coincide with those of the exact operator in (2).

In the present paper we show that the perturbation induced by the approximate operator on the recurrence relation (3) may be far less devastating than the magnitude of the perturbation would indicate. More precisely, the operator \mathbb{A} may be increasingly relaxed, that is ϵ_k may be allowed to increase with k , while the approximation process still converges to a sufficiently accurate approximation of the solution x to the problem (1). We show that this is possible when the components of the approximation in the basis V_m are decreasing in absolute value. We present a new result which helps us quantify this decrease when using the (exact or inexact) FOM and GMRES methods.

* Corresponding Author: e-mail valeria@dm.unibo.it

2 Approximation profile

Write the approximate solution to the problem (1) in $\mathcal{K}_m(\mathbb{A}_\epsilon, v)$ as $x_m = V_m y_m$, for some $y_m \in \mathbb{C}^m$. Then, using (3), the application of the exact operator \mathbb{A} to x_m can be formally written as

$$\mathbb{A}x_m = \mathbb{A}V_m y_m = V_{m+1}H_{m+1,m}y_m + F_{m,\epsilon}y_m.$$

where $V_{m+1}H_{m+1,m}y_m$ corresponds to the application of the relaxed (or perturbed) operator, and $F_{m,\epsilon}y_m$ is the ‘‘correction’’. It thus follows that if $F_{m,\epsilon}y_m$ is small, e.g. in the Euclidean norm, then the application of the unperturbed operator is not significantly affected, since

$$\|\mathbb{A}x_m - V_{m+1}H_{m+1,m}y_m\| = \|F_{m,\epsilon}y_m\|.$$

In other words, if $\|F_{m,\epsilon}y_m\|$ is small, the product $\mathbb{A}x_m$ is well represented by $V_{m+1}H_{m+1,m}y_m$. We explore here when $\|F_{m,\epsilon}y_m\|$ is small. Note that here we are measuring the absolute error in the application of the operator; in the following we implicitly assume that $\|F_{m,\epsilon}y_m\| < \|\mathbb{A}\|$, whereas our analysis shows that $\|F_{m,\epsilon}\|$ does not need to be always smaller than $\|\mathbb{A}\|$. In fact, the original operator may be significantly modified by means of a large $\|F_{m,\epsilon}\|$, while its action on the solution vector is modified by the quantity $\|F_{m,\epsilon}y_m\|$, which in practice may be much smaller.

Indeed, writing the matrix $F_{m,\epsilon}$ as $F_{m,\epsilon} = [f_{\epsilon_1}, f_{\epsilon_2}, \dots, f_{\epsilon_m}]$, and letting η_k , $k = 1, \dots, m$ be the components of y_m , we obtain

$$\|F_{m,\epsilon}y_m\| = \left\| \sum_{k=1}^m f_{\epsilon_k} \eta_k \right\| \leq \sum_{k=1}^m \|f_{\epsilon_k}\| |\eta_k|. \quad (4)$$

Therefore, if for each k , the *product* $\|f_{\epsilon_k}\| |\eta_k|$ is small, then the error in the perturbed application is small. A small product $\|f_{\epsilon_k}\| |\eta_k|$ may be obtained because either of $\|f_{\epsilon_k}\|$ or $|\eta_k|$ is small, and not necessarily both quantities.

In several applications, it was shown that the components of y_m have a decreasing pattern, and this pattern can be related to the residual norm of the problem; we cite some of these below. If we define the problem residual r_k at iteration k as $r_k = G(\mathbb{A}_\epsilon, V_k y_k)$, $k = 1, \dots, m$, then in these applications it can be shown that the j th component of y_m , $\eta_j = e_j^T y_m$, after m iterations satisfies

$$|\eta_j| \leq \gamma_m \|r_k\|, \text{ for } j > k, \quad (5)$$

where γ_m does not depend on k . Assume that an estimate of γ_m is available, where m is the maximum allowed number of iterations, and let ϵ be a small tolerance chosen by the user. Therefore, if at each iteration k the perturbation is required to satisfy $\|f_{\epsilon_k}\| \leq (m\gamma_m \|r_{k-1}\|)^{-1} \epsilon \equiv \epsilon_k$, where $\|r_{k-1}\|$ is the residual norm at the previous iteration, then from (4) we obtain $\|F_{m,\epsilon}y_m\| \leq \epsilon$. From a practical standpoint, this result shows that we can allow for increasing perturbations by a judicious choice of ϵ_k , and still achieve a sufficiently accurate operator evaluation after m iterations.

The bound (5) shows that the components of the solution y_m ‘‘keep memory’’ of the convergence history of the process, in the sense that the components can be bounded by the residual norm at the corresponding (earlier) iteration. It is important to realize that in general, all solution components may change at each iteration, i.e., $e_k^T y_m$ will in general differ from $e_k^T y_{m-1}$. Nonetheless, the process is able to store the convergence information as convergence takes place. Explicit bounds of the form (5) have been shown to hold and used for the following problems of the form (1):

1. Solution of the algebraic linear system $\mathbb{A}x = b$, with $\mathbb{A} = A$; see [8, 10, 1, 5];
2. Solution of the interior eigenvalue problem $Ax = \lambda x$, with $\mathbb{A} = (A - \sigma I)^{-1}$ for some chosen scalar σ [7];
3. Rational function evaluations, which are related to the evaluation of analytic functions such as $\exp(\mathbb{A})b$; Here $\mathbb{A} = (A - \tau I)^{-1}$ for specific choices of the scalar τ ; see [9, 3].

Here we investigate this phenomenon from an approximation viewpoint. Since $\mathcal{K}_m = \text{Range}(V_m) \subset \text{Range}(V_{m+1})$, it may seem reasonable to think that if a good approximation is obtained in $\text{Range}(V_m)$, then increasing the approximation space by one dimension will provide additional information, whose magnitude should be of the order of the current error size. This is typical of iterative refinement procedures. Indeed, in the case of linear systems with methods having decreasing residual norms, the difference $r_m - r_k$ for $k \leq m$ can be bounded as $\|r_m - r_k\| = \|A(V_m y_m - V_k y_k)\| \leq \|r_m\| + \|r_k\| \leq 2\|r_k\|$, so that

$$\|V_m y_m - V_k y_k\| = \|A^{-1}A(V_m y_m - V_k y_k)\| \leq 2\|A^{-1}\| \|r_k\|.$$

Therefore, the refinement in the solution after $m - k$ additional iterations is bounded in terms of the k th residual. In the following proposition we provide an equality between the solutions y_m and y_k . At each step the new approximation is obtained by adding a correction term to the approximate solution of the previous step, whose magnitude is related to the error (or residual) at the previous iteration; note that this fact is well known in case \mathbb{A} is a self-adjoint operator and it is applied exactly. We mention that our result is more general, in that the approximation at one step may be viewed as a correction of

the approximate solution at any previous step. We show this property for the linear system case, with $\mathbb{A} = A$ for two Krylov subspace methods for the solution of $Ax = b$ with $x_0 = 0$: FOM and GMRES, i.e., using Galerkin and Petrov-Galerkin conditions, respectively. To this end, we recall that in FOM, the solution y_m is obtained as the solution to the linear system

$$H_m y_m = e_1 \|b\|, \quad (6)$$

where H_m is the $m \times m$ matrix consisting of the first m rows of $H_{m+1,m}$, while the solution of GMRES is obtained as

$$y_m = \arg \min_{y \in \mathbb{C}^m} \|e_1 \|b\| - H_{m+1,m} y\|; \quad (7)$$

see, e.g., [2], [4], [11]. Note that the following result only depends on the problems (6) and (7) and therefore it holds for both the exact and inexact operators.

Proposition 2.1 *Let $x_k = V_k y_k$ and $x_m = V_m y_m$ be the solutions after k and m iterations with $k < m$, respectively, of either the FOM or GMRES methods. Let $r_k = b - AV_k y_k$ be the corresponding FOM or GMRES residual after k iterations. Then*

$$y_m = \begin{bmatrix} y_k \\ 0 \end{bmatrix} + d_m, \quad \|d_m\| \leq \gamma_m \|r_k\|,$$

where $\gamma_m = \|H_m^{-1}\|$ in the case of FOM and $\gamma_m = \sigma_{\min}(H_{m+1,m})^{-1}$ in the case of GMRES.

Proof. We begin with the proof in the FOM case. Since one has that $\|r_k\| = h_{k+1,k} |e_k^T y_k|$ and H_k is a principal submatrix of H_m , we have that

$$y_m = \begin{bmatrix} y_k \\ 0 \end{bmatrix} - H_m^{-1} \begin{bmatrix} 0 \\ e_1 \end{bmatrix} \|r_k\|,$$

which can be seen to hold by multiplying through by H_m . The result follows.

For the proof in the GMRES case, let $H_{m+1,m} = Q_{m+1} R_{m+1,m}$ be the QR factorization of the rectangular upper Hessenberg matrix $H_{m+1,m}$, where

$$R_{m+1,m} = \begin{bmatrix} R_m \\ 0 \end{bmatrix},$$

and the $m \times m$ matrix R_m is upper triangular. Therefore, $y_m = R_m^{-1} [I_m, 0] Q_{m+1}^T e_1 \|b\|$, where Q_{m+1}^T denotes the complex conjugate of Q_{m+1} . We can write $Q_{m+1}^T = \Omega_{m+1} \Omega_m \cdots \Omega_1$, where for each i , the $(m+1) \times (m+1)$ matrix Ω_i is the Givens rotation applied to $H_{m+1,m}$ to annihilate the element $(i+1, i)$ in $H_{m+1,m}$. The first k rotations $\Omega_1, \dots, \Omega_k$ are also applied to compute y_k , so that

$$y_k = R_k^{-1} [I_k, 0] Q_k^T e_1 \|b\| = R_k^{-1} [I_k, 0] \Omega_k \cdots \Omega_1 e_1 \|b\|. \quad (8)$$

Usually, one has the expression (8) with rotations $\Omega_1, \dots, \Omega_k$ of order $k+1$. Here we are justified in using rotations of order $m+1$ as long as the 0 in $[I_k, 0]$ is the $k \times (m+1-k)$ zero matrix. In addition, notice that $\|r_k\| = |e_{k+1}^T \Omega_k \cdots \Omega_1 e_1| \|b\|$. Let $\widehat{\Omega}_i = [0, I_{k+1:m+1}] \Omega_i [0; I_{k+1:m+1}]$, and since R_k is a principal submatrix of R_m , we write

$$R_m = \begin{bmatrix} R_k & R_{1,2} \\ & R_{2,2} \end{bmatrix}.$$

Then,

$$\begin{aligned} y_m &= R_m^{-1} [I_m, 0] \Omega_m \cdots \Omega_{k+1} \Omega_k \cdots \Omega_1 e_1 \|b\| \\ &= R_m^{-1} [I_m, 0] \Omega_m \cdots \Omega_{k+1} \begin{bmatrix} R_k & \\ & I \end{bmatrix} \begin{bmatrix} R_k^{-1} & \\ & I \end{bmatrix} \Omega_k \cdots \Omega_1 e_1 \|b\| \\ &= R_m^{-1} [I_m, 0] \Omega_m \cdots \Omega_{k+1} \begin{bmatrix} R_k & \\ & I \end{bmatrix} \begin{bmatrix} y_k \\ e_1 \|r_k\| \end{bmatrix} \\ &= R_m^{-1} \begin{bmatrix} R_k & \\ & [I, 0] \widehat{\Omega}_m \cdots \widehat{\Omega}_{k+1} \end{bmatrix} \begin{bmatrix} y_k \\ e_1 \|r_k\| \end{bmatrix} \\ &= \begin{bmatrix} I_k & R_k^{-1} R_{1,2} \\ & R_{2,2} \end{bmatrix}^{-1} \begin{bmatrix} y_k \\ [I, 0] \widehat{\Omega}_m \cdots \widehat{\Omega}_{k+1} e_1 \|r_k\| \end{bmatrix}. \end{aligned}$$

Setting $t = [I, 0] \widehat{\Omega}_m \cdots \widehat{\Omega}_{k+1} e_1 \|r_k\|$, with $\|t\| \leq \|r_k\|$, it follows that

$$\begin{aligned} y_m &= \begin{bmatrix} y_k - R_k^{-1} R_{1,2} R_{2,2}^{-1} t \\ R_{2,2}^{-1} t \end{bmatrix} \\ &= \begin{bmatrix} y_k \\ 0 \end{bmatrix} + \begin{bmatrix} I & -R_k^{-1} R_{1,2} R_{2,2}^{-1} \\ & R_{2,2}^{-1} \end{bmatrix} \begin{bmatrix} 0 \\ t \end{bmatrix} = \begin{bmatrix} y_k \\ 0 \end{bmatrix} + \begin{bmatrix} I & R_k^{-1} R_{1,2} \\ & R_{2,2} \end{bmatrix}^{-1} \begin{bmatrix} 0 \\ t \end{bmatrix} \\ &= \begin{bmatrix} y_k \\ 0 \end{bmatrix} + \begin{bmatrix} R_k & R_{1,2} \\ & R_{2,2} \end{bmatrix}^{-1} \begin{bmatrix} R_k & \\ & I \end{bmatrix} \begin{bmatrix} 0 \\ t \end{bmatrix} = \begin{bmatrix} y_k \\ 0 \end{bmatrix} + R_m^{-1} \begin{bmatrix} 0 \\ t \end{bmatrix}, \end{aligned}$$

from which the result follows, recalling that $\|R_m^{-1}\| = \sigma_{\min}(H_{m+1,m})^{-1}$. \square

The result of Proposition 2.1 shows an explicit relation of each *single* component of the two vectors. In the case of GMRES, from the proof we obtain $[I_k, 0]y_m = y_k - R_k^{-1} R_{1,2} R_{2,2}^{-1} t$, so that for $j \leq k$,

$$e_j^T y_m = e_j^T y_k - e_j^T R_k^{-1} R_{1,2} R_{2,2}^{-1} t. \quad (9)$$

Since R_k^{-1} is upper triangular, only the last $k - j$ rows of $R_{1,2}$ are involved in the computation of (9). Unless R_m^{-1} has large entries, the first k components of y_m stabilize at their final value if $\|r_k\|$ is sufficiently small. Note that the result of Proposition 2.1 gives a different proof of the result shown in [8, 10], while providing more insight into the way the approximation is improved as the Krylov subspace dimension grows. For this reason, we believe the relation of Proposition 2.1 may give further insight in the understanding of truncation strategies in Krylov subspace methods. It is also worth noticing that similar relations hold for the restarted versions of the methods above; see [6].

We remark that an update recurrence of the type shown in Proposition 2.1 has been given in [7] for approximate eigenvectors of $Ax = \lambda x$ projected in the approximation subspace $\mathcal{K}_m(\mathbb{A}_\epsilon, v)$; cf. [7, Proposition 2.2].

A similar result can also be obtained when dealing with linear systems whose coefficient matrix is a polynomial in A with distinct roots. In this case, by using the partial fraction expansion of the inverse of the polynomial, the solution can be recast as the solution of single linear systems in A . More precisely, assume that $\Phi_k(A)x = b$ is the linear system to be solved, with $\Phi_k(\lambda) = \prod_{j=1}^k (\lambda - \lambda_j)$. Using $\Phi_k(\lambda)^{-1} = \tau_0 + \sum_{j=1}^k \tau_j (\lambda - \lambda_j)^{-1}$, for some scalars τ_0, \dots, τ_k , we can write

$$x = \Phi_k(A)^{-1} b = \tau_0 b + \sum_{j=1}^k \tau_j (A - \lambda_j I)^{-1} b.$$

Therefore, an approximation $x_m = V_m y_m$ to x may be obtained by solving the k shifted systems in A in the same Krylov subspace generated by A ; cf. [3] and references therein. By applying the result of Proposition 2.1 to each linear system, a decreasing pattern for y_m can be readily obtained, where the components decrease according to the residual history of the slowest converging system.

Acknowledgements The second author was supported in part by the U.S. National Science Foundation under grant DMS-0207525.

References

- [1] A. BOURAS AND V. FRAYSSÉ, *Inexact matrix-vector products in Krylov methods for solving linear systems: A relaxation strategy*, SIAM J. Matrix Anal. Appl., 26 (2005), pp. 660–678.
- [2] A. GREENBAUM, *Iterative Methods for Solving Linear Systems*. SIAM, Philadelphia, 1997.
- [3] L. LOPEZ AND V. SIMONCINI, *Analysis of projection-type methods for approximating the matrix exponential operator*, tech. rep., 2005. Submitted.
- [4] Y. SAAD, *Iterative methods for sparse linear systems*. Second edition, SIAM, Philadelphia, 2003.
- [5] M. SARKIS AND D. SZYLD, *A Proposal for a dynamically adapted inexact additive Schwarz preconditioner*. Proceedings of the Sixteenth Conference on Domain Decomposition Methods, 2005. Lecture Notes in Computational Science and Engineering, Springer, to appear.
- [6] V. SIMONCINI, *On the convergence of restarted Krylov subspace methods*, SIAM J. Matrix Analysis and Appl., 22 (2000), pp. 430–452.
- [7] V. SIMONCINI, *Variable accuracy of matrix-vector products in projection methods for eigencomputation*, tech. rep., 2004. SIAM J. Numer. Analysis, to appear.
- [8] V. SIMONCINI AND D. B. SZYLD, *Theory of inexact Krylov subspace methods and applications to scientific computing*, SIAM J. Sci. Comput., 25 (2003), pp. 454–477.
- [9] J. VAN DEN ESHOF AND M. HOCHBRUCK, *Preconditioning Lanczos approximations to the matrix exponential*, tech. rep., Mathematical Institute, Heinrich Heine Universität, Düsseldorf, March 2004. SIAM J. Scient. Computing, to appear.
- [10] J. VAN DEN ESHOF AND G. SLEIJPEN, *Inexact Krylov subspace methods for linear systems*, SIAM J. Matrix Anal. Appl., 26 (2004), pp. 125–153.
- [11] H. A. VAN DER VORST, *Iterative Krylov Methods for Large Linear Systems*. Cambridge University Press, Cambridge, 2003.