Minimal residual methods for large scale Lyapunov equations

Yiding Lin [a,b], Valeria Simoncini [b,*]

[a] *School of Mathematical Sciences, Xiamen University, China*
[b] *Dipartimento di Matematica, Università di Bologna, Piazza di Porta S. Donato 5, 40127 Bologna, Italy*

A B S T R A C T

Projection methods have emerged as competitive techniques for solving large scale matrix Lyapunov equations. We explore the numerical solution of this class of linear matrix equations when a Minimal Residual (MR) condition is used during the projection step. We derive both a new direct method, and a preconditioned operator-oriented iterative solver based on CGLS, for solving the projected reduced least squares problem. Numerical experiments with benchmark problems show the effectiveness of an MR approach over a Galerkin procedure using the same approximation space.

© 2013 IMACS. Published by Elsevier B.V. All rights reserved.

1. Introduction

The solution of the following Lyapunov matrix equation,

$$AX + XA^* + BB^* = 0, \quad A \in \mathbb{R}^{n \times n}, B \in \mathbb{R}^{n \times p} \tag{1.1}$$

with $p \ll n$ and $*$ denoting the conjugate transpose, plays an important role in the stability analysis of dynamical systems and in control theory; see, e.g., [1,4]. We shall assume that A is stable, that is the real parts of its eigenvalues are in the left half complex plane. For n small, say up to a few hundreds, well established and robust numerical methods exist, which are based on the Schur decomposition of A [2]; see also [16]. For large scale problems, and in particular for A stemming from the discretization of two- (2D) or three-dimensional (3D) partial differential operators, several numerical methods have been explored to approximate the positive semidefinite solution matrix X by means of a low-rank matrix, $ZZ^* \approx X$, so that only the tall matrix Z needs to be stored. The Alternating Direction Implicit (ADI) method was introduced by Wachspress [43] as an iterative procedure to obtain a full approximation matrix; see also [33]. An effective cyclic low-rank version was proposed by Penzl in [37] and by Li and White in [32], and since then, a lot of work has been devoted to the algorithmic and convergence analysis of the method; we refer to [6] for a recent overview, and to [26,5] for extended ADI-based strategies.

Projection methods have also been proposed for iteratively solving (1.1). As for modern ADI schemes, these methods provide a factorized approximate solution, obtained from the projection of the original problem onto a much smaller dimension space. A pioneering strategy was proposed by Saad in 1990 [39] where a Krylov subspace was used as projection space, and then further explored in [22,25,13]; see also [27]. However, only more recently projection strategies have emerged as true competitors with respect to ADI type methods: the use of powerful subspaces for the projection step may allow one to determine an accurate factored approximation with a very small approximation space dimension. In particular, extended Krylov subspaces and more general rational Krylov subspaces have proven to be particularly effective on large 2D and 3D

* Corresponding author at: Dipartimento di Matematica, Università di Bologna, Piazza di Porta S. Donato 5, 40127 Bologna, Italy.
E-mail addresses: Yiding.Lin@gmail.com (Y. Lin), valeria.simoncini@unibo.it (V. Simoncini).

problems [41,11]. Due to the good experimental behavior, large efforts have been recently devoted to the theoretical analysis of projection methods, whose new results have filled a gap in the understanding of these strategies for matrix equations, also with respect to the related ADI method [12,42,3,31].

Projection-type methods for (1.1) generally perform the following steps:

1. Generate the approximation subspace Range($\mathcal{V}_k$).
2. Project the original problem onto Range($\mathcal{V}_k$). solve the reduced problem for Y_k
3. Generate the approximate solution $X_k = \mathcal{V}_k Y_k \mathcal{V}_k^*$.

If the approximate solution is not satisfactory, then the space is expanded and a new candidate computed. The second step above characterizes the specific extraction method, once the approximation space has been chosen. Most successful projection methods use the Galerkin condition for extracting the approximate solution from the given space. This condition requires that the associated residual $R = AX_k + X_k A^* + BB^*$ be orthogonal to the approximation space. In matrix terms, this translates into the equation (cf., e.g., [39])

$$\mathcal{V}_k^*\big(A\mathcal{V}_k Y_k \mathcal{V}_k^* + \mathcal{V}_k Y_k \mathcal{V}_k^* A^* + BB^*\big)\mathcal{V}_k = 0. \tag{1.2}$$

Assuming that $\mathcal{V}_k$ has orthonormal columns, Y_k can be determined as the solution to a reduced Lyapunov matrix equation with coefficient matrix $\mathcal{V}_k^* A \mathcal{V}_k$.

We are interested in exploring a different (*Petrov–Galerkin*-type) condition, which corresponds to minimizing the Frobenius norm of the residual. In this case, the reference problem is

$$Y_k^{MR} = \arg \min_{Y_k \in \mathbb{R}^{k \times k}} \big\| A\mathcal{V}_k Y_k \mathcal{V}_k^* + \mathcal{V}_k Y_k \mathcal{V}_k^* A^* + BB^* \big\|_F. \tag{1.3}$$

The approximation process can thus be stopped by monitoring the quantity that is actually minimized. On the contrary, the residual norm in the Galerkin process may exhibit a erratic behavior, thus delaying the process. This behavior is reminiscent of the well-known situation in projection-based methods for solving standard linear systems [40]: the residual norm of the FOM method (applying the Galerkin condition) may heavily oscillate, whereas the minimized residual norm of GMRES (based on a residual minimization procedure) smoothly decays. A thorough analysis shows that peaks and plateaux of FOM and GMRES, respectively, can be explicitly related [9,8]. In the matrix equation setting we can numerically confirm this fascinating interplay.

Minimal residual approaches were analyzed in [25,22]; however the proposed algorithms for solving the reduced problem did not lead to computationally competitive methods with respect to the Galerkin strategy. As a result, minimal residual strategies were somehow disregarded, in spite of their good theoretical properties. In this paper we provide an implementation of the minimal residual method (MR in the following) whose leading computational cost is comparable to that of the Galerkin method, using the same projection space, with smoother convergence behavior. The algorithmic bottleneck in the MR method is the solution of a reduced order linear least squares matrix problem with three terms, whose efficient numerical solution has not been addressed in the standard least squares literature. This cost should be compared with that of solving the reduced Lyapunov equation (1.2) for the Galerkin approach. By using Kronecker products, the least squares problem can be reformulated as a standard least squares problem of exploding size, so that solving the latter problem becomes prohibitive even for moderate dimension of the matrix least squares problem. We thus explore three venues: (i) We propose a new direct method that computes the residual norm and the least squares solution by using spectral decompositions; (ii) We improve the algorithm originally proposed in [22] by exploiting the Lyapunov structure of our setting; (iii) We devise operator oriented preconditioning strategies for iteratively solving the matrix normal equation associated with the matrix least squares problem. For all strategies we provide a complete derivation and discuss the main algebraic properties and computational challenges. A selection of numerical experiments will be reported to show the potential of the new algorithms on various known data sets.

The paper is organized as follows: given an approximation space satisfying certain properties, Section 2 describes the projection step under a minimal residual constraint. Section 3 serves as introduction to the description of the computational core. In Section 3.1 we devise a procedure for computing the least squares residual norm as the space expands, without explicitly computing the least squares solution; from Section 3.2 to Section 3.4 we propose three different strategies, direct and iterative, to compute the final solution to the reduced problem. Section 4 summarizes the minimal residual algorithm. In Section 5 we report on our numerical experience with these algorithms when different approximation spaces are used, and compare their performance with that of a standard Galerkin approach. Our conclusions and open problems are described in Section 6.

Throughout the paper, $\|\cdot\|_F$ denotes the Frobenius norm, the symbol $\otimes$ denotes the Kronecker product and vec($\cdot$) denotes the vectorization operator, which stacks all columns of a matrix one below the other. The field of values of A is defined as $F(A) = \{x^* A x: \ x \in \mathbb{C}^n, x^* x = 1\}$. The identity matrix of size k is denoted by I_k, though the subscript will be omitted when clear from the context. E_k collects the last k columns of the identity matrix, whose dimension is clear from the context. Whenever possible, Matlab [34] notation will be used, e.g., $[x; y] = [x^*, y^*]^*$. Notation corresponding to Matlab functions, such as diag, blkdiag, etc., will also be used for convenience.