

COMPUTATIONAL METHODS FOR LINEAR MATRIX EQUATIONS*

V. SIMONCINI[†]

Abstract. Given the square matrices A, B, D, E and the matrix C of conforming dimensions, we consider the linear matrix equation $AXE + DXB = C$ in the unknown matrix X . Our aim is to provide an overview of the major algorithmic developments that have taken place in the past few decades in the numerical solution of this and of related problems, which are becoming a reliable numerical tool in the formulation and solution of advanced mathematical models in engineering and scientific computing.

Key words. Sylvester equation, Lyapunov equation, Stein equation, multiple right-hand side, generalized matrix equations. Schur decomposition. Large scale computation.

AMS subject classifications. 65F10, 65F30, 15A06

Contents

1.	Introduction	2
2.	Notation and preliminary definitions	4
3.	Applications	5
4.	Continuous-time Sylvester equation	10
4.1.	Stability of the Sylvester equation	12
4.2.	Small scale computation	13
4.3.	Computation with large A and small B	15
4.4.	Computation with large A and large B	19
4.4.1.	Projection methods	21
4.4.2.	ADI iteration	25
4.4.3.	Data sparse and other methods	27
5.	Continuous-time Lyapunov equation	28
5.1.	Small scale computation	29
5.2.	Large scale computation	31
5.2.1.	Projection methods	31
5.2.2.	ADI methods	35
5.2.3.	Spectral, Sparse format and other methods	39
6.	The discrete Sylvester and Stein equations	42
7.	Generalized linear equations	44
7.1.	The Generalized Sylvester and Lyapunov equations	44
7.2.	Bilinear, constrained, and other linear equations	46
7.3.	Sylvester-like and Lyapunov-like equations	50
8.	Software and high performance computation	52
9.	Concluding remarks and future outlook	53
10.	Acknowledgements	54

*Version of January 21, 2014.

[†]Dipartimento di Matematica, Università di Bologna, Piazza di Porta San Donato 5, I-40127 Bologna, Italy (valeria.simoncini@unibo.it).

1. Introduction. Given the real or complex square matrices A, D, E, B and the matrix C of conforming dimensions, we consider the linear matrix equation

$$AXE + DXB = C \quad (1.1)$$

in the unknown matrix¹ \mathbf{X} , and its various generalizations. If E and D are identity matrices, then (1.1) is called the Sylvester equation, as its first appearance is usually associated with the work of J.J. Sylvester [238]; if in addition $B = A^*$, where A^* is the conjugate transpose of A , then the equation is called the Lyapunov equation in honor of A. M. Lyapunov and his early contributions to the stability problem of motion; see [14] and the whole issue of the same journal. We shall mainly consider the generic case, thus assuming that all the involved matrices are nonzero.

Under certain conditions on the coefficient matrices, (1.1) has a unique solution, with available elegant and explicit closed forms. These are usually inappropriate as computational devices, either because they involve estimations of integrals, or because their computation can be polluted with numerical instabilities of various sorts. Nonetheless, closed forms and other properties of the solution matrix have strongly influenced the computational strategies that have led to most algorithms used today for numerically solving (1.1), in the case of small or large dimensions of the coefficient matrices. Due to the availability of robust and reliable core algorithms, (1.1) now arises in an increasingly larger number of scientific computations, from statistics to dynamical systems analysis - with a major role in control applications, and also as a workhorse of more computationally intensive methods. In section 3 we will briefly review this broad range of numerical and application problems.

Our aim is to provide an overview of the major algorithmic developments that have taken place in the past few decades in the numerical solution of (1.1) and of related problems, both in the small and large scale cases. A distinctive feature in the large scale setting is that although the coefficient matrices may be sparse, the solution matrix is usually dense and thus impossible to store in memory. Therefore, ad-hoc strategies need to be devised to approximate the exact solution in an affordable manner.

Functions related to the solution matrix \mathbf{X} such as the spectrum, the trace and the determinant, also have an important role in stability analysis and other applications. Although we shall not discuss in detail the computational aspects associated with these functions, we shall occasionally point to relevant results and appropriate references.

Linear matrix equations have received considerable attention since the early 1900's, and have been the topic of many elegant and thorough studies in the 1950's and 1960's, which use deep tools of matrix theory and functional analysis. The field continues to prosper with the analysis of new challenging extensions of the main equation (1.1), very often stimulated by application problems. Our contribution is intended to focus on the computational methods for solving these equations. For this reason, in our presentation we will mostly sacrifice the theoretical results, for which we refer the interested reader to, e.g., [88], [166], [130], [16].

The literature on the Lyapunov equation is particularly rich, due to the prominent role of this matrix equation in control. In particular, many authors have focused on numerical strategies associated specifically to this equation. As a consequence, the Sylvester and Lyapunov equations have somehow evolved differently. For these

¹Here and in the following we shall use bold face letters to denote the unknown solution matrices.

reasons, and to account for the literature in a homogeneous way, we shall first discuss numerical strategies for the Sylvester equation, and then treat in detail the Lyapunov problem. For A and B of size up to a few hundreds, the Schur-decomposition based algorithm by Bartels and Stewart ([15]) has since its appearance become the main numerical solution tool. In the large scale case, various directions have been taken, and a selection of effective algorithms is available, from projection methods to sparse format iterations. Despite a lot of intense work in the past 15-20 years, the community has not entirely agreed upon the best approaches for all settings; hence the need for an overview that aims at analyzing where the field stands at this point.

For A and B of the order of 10^4 or larger, the solution \mathbf{X} cannot be store explicitly; current memory effective strategies rely on factored low rank or sparse approximations. The possibility of computing a memory conserving good approximate solution in the large scale case highly depends on the data. In particular, for C definite, accurate low-rank approximations may be hard, if not impossible, to find. For instance, the equation $A\mathbf{X} + \mathbf{X}A^\top = I$ with A nonsingular and symmetric admits the unique solution $\mathbf{X} = \frac{1}{2}A^{-1}$, which is obviously full rank, with not necessarily quickly decreasing eigenvalues, so that a good low rank approximation cannot be determined.

The distinction between small, moderate and large size is clearly architecture dependent. In the following we shall refer to “small” and medium problem size when the coefficient matrices have dimensions of a few thousands at most; on high performance computers these digits can be considerably larger. Small and medium size linear equations can be solved with decomposition-based methods on laptops with moderate computational efforts. The target for current large-scale research are matrices of dimensions $\mathcal{O}(10^6)$ or larger, with a variety of sparsity patterns.

Throughout the paper we shall assume that E, D are either the identity, or that at least one of them is nonsingular. Singular E, D have great relevance in control applications associated with differential-algebraic equations and descriptor systems but require a specialized treatment, which can be found, for instance, in [162].

Equation (1.1) is a particular case of the linear matrix equation

$$A_1\mathbf{X}B_1 + A_2\mathbf{X}B_2 + \dots + A_k\mathbf{X}B_k = C, \quad (1.2)$$

with A_i, B_i , $i = 1, \dots, k$ square matrices, and C of dimension $n \times m$. While up to 15-20 years ago this multi-term equation could be rightly considered of mainly theoretical interest, the recent developments associated with problems stemming from applications with parameters or a dominant stochastic component have brought multi-term linear matrix equations to play a fundamental role; see section 3 and section 7.2 for applications and references. Equation (1.2) is very difficult to analyze in its full generality, and necessary and sufficient conditions for the existence and uniqueness of the solution \mathbf{X} explicitly based on $\{A_i\}, \{B_i\}$, are hard to get, except for some very special cases [166],[155]. While from a theoretical view point the importance of taking into account the structure of the problem has been acknowledged [155], this has not been so for computational strategies, especially for large scale problems. The algorithmic device most commonly used for (1.2) consists in transforming the matrix equation above into a vector form by means of the Kronecker product (defined below). The problem of the efficient numerical solution of (1.2), with a target complexity of at most $\mathcal{O}(n^3 + m^3)$ operations, has only recently started to be addressed. The need for a low complexity method is particularly compelling whenever either or both A_i and B_i have large dimensions. Approaches based on the Kronecker formulations were abandoned for (1.1) as core methods, since algorithms with a complexity of a modest

power of the coefficient matrices dimension are now available. The efficient numerical solution to (1.2) thus represents the next frontier for linear matrix equations, so as to assist rapidly developing application models.

Various forms of generalizations have also been tackled in the literature, as they are more and more often encountered in applications. This is the case, for instance, for bilinear equations (in two unknown matrices), and for *systems* of bilinear equations. These are an open computational challenge, especially in the large scale case, and their efficient numerical solution would provide a great advantage for emerging mathematical models; we discuss these generalizations in section 7.

A very common situation arises when $B = 0$ and C is tall in (1.1), so that the matrix equation reduces to a standard linear system with multiple right-hand sides, the columns of C . This is an important problem in application, and a significant body of literature is available, with a vast number of contributions in the past twenty years. Since the most popular procedures for the solution of $A\mathbf{X} = C$ are usually derived from the single right-hand side case, we shall not discuss them here, as the topic surely deserves a dedicated treatment; instead, we refer the reader to [212] and to the more recent list of references in [111].

After a brief account in section 3 of the numerous application problems where linear matrix equations arise, we shall recall the main properties of these equations, together with possible explicit forms for their solution matrix. The rest of this paper describes many approaches that have been proposed in the recent literature: we first treat the Sylvester equation, when A and B are small, when one of the two is large, and when both are large. Indeed, rather different approaches can be employed depending on the size of the two matrices. We shall then focus on the Lyapunov equation: due to its relevance in control, many developments have specifically focused on this equation, therefore the problem deserves a separate treatment. We shall describe the algorithms that were specifically designed to take advantage of the symmetry, while we shall only mention the solution methods that are common to the Sylvester equation. The small-scale problem is computationally well understood, whereas the large-scale case has seen quite significant developments in the past ten years. Later sections report on the computational devices associated with the numerical solution of various generalizations of (1.1), which have been taking place in the past few years.

2. Notation and preliminary definitions. Unless stated otherwise, throughout the paper we shall assume that the coefficient matrices are real. Moreover, $\text{spec}(A)$ denotes the set of eigenvalues of A , and A^\top , A^* denote that transpose and conjugate transpose of A , respectively. For $z \in \mathbb{C}$, \bar{z} is the complex conjugate of z .

A matrix A is *stable* if all its eigenvalues have negative real part, and *negative definite* if for all unit 2-norm complex vectors x , the quantity x^*Ax has negative real part, namely, the field of values $W(A) = \{z \in \mathbb{C} : z = x^*Ax, x \in \mathbb{C}^n, \|x\| = 1\}$ is all contained in the open left half complex plane. The notation $A \succ 0$ ($A \succeq 0$) states that A is a Hermitian and positive definite (semi-definite) matrix.

The vector e_i denotes the i th column of the identity matrix, whose dimension will be clear from the context; I_n denotes the identity matrix of size n , and the subscript will be omitted when clear from the context. Throughout, given $x \in \mathbb{C}^n$, $\|x\|$ denotes the 2-norm of x , $\|A\|$ or $\|A\|_2$ denotes the matrix norm induced by the vector 2-norm, while $\|A\|_F$ denotes the Frobenius norm of $A = (a_{i,j})_{i=1,\dots,n,j=1,\dots,m}$, that is $\|A\|_F^2 = \sum_{i,j} |a_{i,j}|^2$. The notation $[A; B]$ will be often used to express the matrix obtained by stacking the matrix B below the matrix A , both having conforming dimensions.

For given matrices $A \in \mathbb{C}^{n_A \times m_A}$, $A = (a_{ij})_{i=1, \dots, n_A, j=1, \dots, m_A}$ and $B \in \mathbb{C}^{n_B \times m_B}$, the Kronecker product is defined as

$$A \otimes B = \begin{bmatrix} a_{11}B & a_{12}B & \cdots & a_{1m_A}B \\ a_{21}B & a_{22}B & \cdots & a_{2m_A}B \\ \vdots & \vdots & \ddots & \vdots \\ a_{n_A1}B & a_{n_A2}B & \cdots & a_{n_Am_A}B \end{bmatrix} \in \mathbb{C}^{n_A n_B \times m_A m_B}; \quad (2.1)$$

the vec operator stacks the columns of a matrix $X = [x_1, \dots, x_m] \in \mathbb{C}^{n \times m}$ one after the other as

$$\text{vec}(X) = \begin{bmatrix} x_1 \\ \vdots \\ x_m \end{bmatrix} \in \mathbb{C}^{nm \times 1}.$$

We summarize some well known properties of the Kronecker product in the next lemma; see, e.g., [130].

LEMMA 2.1. *Some properties:*

- (i) $\text{vec}(AXB) = (B^\top \otimes A)\text{vec}(X)$ (note the complex transposition for B);
- (ii) If $A \in \mathbb{R}^{n \times n}$ and $B \in \mathbb{R}^{m \times m}$, and $\lambda_A \in \text{spec}(A)$, $\lambda_B \in \text{spec}(B)$, then $\lambda_A \lambda_B \in \text{spec}(A \otimes B)$; (and every eigenvalue of $A \otimes B$ is the product of eigenvalues of A and B)
- (iii) Under the hypotheses of (ii), $\lambda_A + \lambda_B \in \text{spec}(I_m \otimes A + B \otimes I_n)$; (and every eigenvalue of $I_m \otimes A + B \otimes I_n$ is the sum of eigenvalues of A and B .)

3. Applications. Matrix equations are ubiquitous in signal processing, control and system theory; see, e.g., [4],[248],[88],[66],[26],[216],[59] and references therein. Most time-dependent models may be represented as linear or non-linear dynamical systems, accounting for the prediction, simulation and control of real world phenomena. The numerical solution of matrix equations and the relevance of its role within engineering applications justify the great effort put into this problem by the scientific community, both from a mathematical and from an applied angle. Special issues of journals and multi-contribution books are often devoted to advances in this and related areas, attesting the search for new algorithms that can take into account the properties of the problem, such as structure, size and functional characterizations.

Linear matrix equations have an important role in the stability analysis of linear dynamical systems, and take also part in the theoretical developments of non-linear ones. Consider the following continuous-time linear system²

$$\dot{x} = Ax + B_1 u, \quad y = B_2^\top x, \quad (3.1)$$

where x is the model state, u is the input and y is the output, and the matrices A , B_1 and B_2 are time-invariant. Assuming A is stable, that is its eigenvalues have negative real part, then the solutions \mathbf{P} and \mathbf{Q} to the following Lyapunov equations

$$A\mathbf{P} + \mathbf{P}A^\top + B_1 B_1^\top = 0, \quad A^\top \mathbf{Q} + \mathbf{Q}A + B_2 B_2^\top = 0,$$

are called the controllability and observability Gramians, respectively, and they are used, for instance, to measure the energy transfers in the system (3.1)[4, sec.4.3.1].

²In the control literature, B_1, B_2 are usually denoted by B and C^\top , respectively; we opted for a slightly different notation because here B and C have a different meaning.

Under certain additional hypotheses it may be shown that the symmetric matrices \mathbf{P} and \mathbf{Q} are positive definite. These two latter matrices are key when one is interested in reducing the original system into one of much smaller dimension, while essentially preserving the main dynamical system properties. Indeed *balanced reduction*, which was originally used to improve the sensitivity to round-off propagation in filter design [186], determines an appropriate representation basis for the system such that the Gramians are equal and diagonal [4], so that the reduction of that basis will maintain this property of the Gramians. The diagonal Gramians then contain information on the output error induced by the reduced model.

Alternatively, if B_1 and B_2 have the same number of columns, one can solve the following Sylvester equation,

$$A\mathbf{W} + \mathbf{W}A + B_1B_2^\top = 0,$$

thus obtaining the *cross-Gramian* \mathbf{W} [84], which contains information on controllability and observability of the system. For B_1, B_2 having a single column, or for A symmetric and B_1, B_2 such that $B_2^\top(zI - A)^{-1}B_1$ is symmetric, it is possible to show that $\mathbf{W}^2 = \mathbf{P}\mathbf{Q}$, so that the eigenvalues of \mathbf{W} coincide with the square root of the eigenvalues of $\mathbf{P}\mathbf{Q}$ [85],[234]. In general, the latter are called the Hankel singular values of the system, and they are invariant under state space transformations; see [4] for a detailed discussion of these quantities and their role in model order reduction. A different Sylvester equation was used in [89] to derive a numerical algorithm that couples the two Gramians \mathbf{P} and \mathbf{Q} . Similar results can be stated for the case of the discrete-time time-invariant linear systems

$$\begin{aligned} x(k+1) &= Ax(k) + B_1u(k) \\ y(k) &= B_2^\top x(k), \end{aligned} \tag{3.2}$$

which are associated, for instance, with the discrete-time Lyapunov equation

$$A\mathbf{X}A^\top - \mathbf{X} + B_1B_1^\top = 0.$$

As a particular case of the linear equation in (1.1), the generalized Lyapunov equation

$$A\mathbf{X}E^\top + E\mathbf{X}A^\top = C, \tag{3.3}$$

has a special interest in control; see also recent applications in Hopf bifurcation identification in linear stability analysis [80],[182]. The case $E \neq I$ arises in a control problem, for instance, whenever a second or higher order ordinary differential equation is discretized. Consider the linear time-invariant second-order system

$$\begin{aligned} Mq''(t) + Dq'(t) + Kq(t) &= B_2u(t) \\ C_2q'(t) + C_1q(t) &= y(t), \end{aligned}$$

where $q(t) \in \mathbb{R}^n$ is the displacement, and $u(t) \in \mathbb{R}^m$, $y(t) \in \mathbb{R}^p$ are the control input and output, respectively. Then by defining the matrices

$$E = \begin{bmatrix} I & \\ & M \end{bmatrix}, \quad A = \begin{bmatrix} 0 & I \\ -K & -D \end{bmatrix}, \quad B = \begin{bmatrix} 0 \\ B_2 \end{bmatrix} \quad \text{and} \quad C = [C_1, C_2],$$

the second-order system can be rewritten as a first order linear system

$$Ex'(t) = Ax(t) + Bu(t), \quad y(t) = Cx(t),$$

with $x(t) = [q(t); q'(t)]$, whose stability analysis gives rise to (3.3).

The Sylvester equation is classically employed for the design of Luenberger observers [180]; we refer the reader to section 7 for a more detailed discussion. Linear matrix equations are also used in control as a technical tool for solving other problems, see, e.g., [269],[88],[168],[191], and for the reduction of nonlinear models; see, e.g., [216],[26] and references therein.

The Sylvester equation often occurs in linear and generalized eigenvalue problems for the computation of invariant subspaces by means of the Riccati equation [236], [227], [68]. In fact, the algebraic Riccati equation itself, defined in the symmetric case as

$$A^\top X + XA - XFX + G = 0, \quad (3.4)$$

with F and G symmetric, provides a formidable setting for linear matrix equations: this quadratic equation is sometimes dealt with by solving a sequence of linear Sylvester or Lyapunov equations with possibly varying known term and coefficient matrices. The following Newton-Kleinman iteration is one of the leading methods for solving (3.4) in the large scale case, whenever $F = BB^\top$ and $G = C^\top C$ have low rank:

Algorithm 1. Given $X_0 \in \mathbb{R}^{n \times n}$ such that $X_0 = X_0^\top$, $A^\top - X_0BB^\top$ is stable

1. **For** $k = 0, 1, \dots$, until convergence
2. **Set** $\mathcal{A}_k^\top = A^\top - X_kBB^\top$
3. **Set** $\mathcal{C}_k^\top = [X_kB, C^\top]$
4. **Solve** $\mathcal{A}_k^\top X_{k+1} + X_{k+1}\mathcal{A}_k + \mathcal{C}_k^\top \mathcal{C}_k = 0$

At each iteration the most computationally intensive operation is step 4, which requires the solution of a Lyapunov equation, whose data changes at each iteration [42].

With the aim of controlling resonance modes in vibrating structures, Sylvester equations also arise in solving quadratic eigenvalue assignment problems, see, e.g., [48]. Large eigenvalue problems are also a key step in the detection of a Hopf bifurcation in large-scale dynamical systems that depend on some physical parameters. However, it is possible to compute these parameters without actually computing the relevant eigenvalues. In [182], it was shown that this can be performed by means of a *matrix* inverse iteration procedure, which involves approximately solving a sequence of large-scale Lyapunov equations; see also [80]. Lyapunov equations are a theoretical and computational tool also in hydrodynamic stability theory of time-dependent problems, which is emerging as an attractive alternative to classical modal analysis, in the quantitative description of short-term disturbance behaviors [217]. A large list of references on application problems where the Lyapunov equation plays an important role is available in the last chapter of [88].

Different application areas have emerged that can take advantage of an efficient solution of linear matrix equations. Problems associated with image processing seem to provide a rich source. For instance, Sylvester equations can be used to formulate the problem of restoration of images affected by noise [53]. The degraded image can be written as $\mathbf{g} = \mathbf{f} + \eta$, where η is the Gaussian noise vector. A linear operator (filter) L is applied to \mathbf{g} to determine an estimate $\hat{\mathbf{f}} := L\mathbf{g}$ of the original image. A possible choice for L is the Wiener filter $L = \Phi_f(\Phi_f + \Phi_\eta)^{-1}$, where Φ_η is the covariance matrix of the noise, while $\Phi_f = \Phi_y \otimes \Phi_x$ is the covariance of f , assuming that the variability in the vertical (y) and horizontal (x) directions are unrelated.

The minimum mean square error estimate $\hat{\mathbf{f}}$ of \mathbf{f} can be computed by solving the linear system $(I + \Phi_\eta \Phi_f^{-1})\hat{\mathbf{f}} = \mathbf{g}$. For $\Phi_\eta = \sigma_\eta^2 I$, corresponding to white noise η and Gaussian with variance σ_η^2 , the system is given by

$$(I + \sigma_\eta^2 \Phi_y^{-1} \otimes \Phi_x^{-1})\hat{\mathbf{f}} = \mathbf{g},$$

which is nothing but the Kronecker formulation of a Sylvester equation.

A similar optimization model can be used in adaptive optics, a technology developed for compensation of aberrations in optical systems or due to atmospheric turbulence, mainly used in high quality astronomical observations and measurements [206]. Within the image processing application, the problem of estimating a 3D object's pose obtained from 2D image sequences can be stated as a constrained optimization problem [57]. This leads to the solution of a sequence of small Sylvester equations. In fact, depending on the number of poses, the occurring linear matrix equations have more than two terms, and can be formulated as in (1.2); see [57].

The Sylvester equation was highlighted as a *model problem* in the solution of elliptic boundary value problems governed by the two-dimensional differential operator

$$\mathcal{L}(u) = -\nabla \cdot (\kappa \nabla u)$$

by Ellner and Wachspress [78]: they devised a matrix algorithmic version of the (differential) ADI algorithm by Peaceman and Rachford, and this became the foundation of ADI-type methods for linear matrix equations. Wachspress showed that the constant coefficient second order differential equation can be used as preconditioner for the original operator, and that the application of the preconditioner amounts to solving a Lyapunov equation [256]. Sylvester equations can also be used in the implementation of implicit Runge-Kutta integration formulae and block multi-step formulae for the numerical solution of ordinary differential equations [82].

Discrete-time Sylvester and Lyapunov equations (see section 6) also arise for instance in statistics and probability ([151],[150],[149],[10]), and as a building block for solving the discrete-time algebraic Riccati equation [42].

Similarly to the Sylvester equation, the multi-term matrix equation (1.2) may be viewed as a model problem for certain convection-diffusion partial differential equations. For instance, let us consider the following two-dimensional problem with *separable* coefficients:

$$-\varepsilon u_{xx} - \varepsilon u_{yy} + \phi_1(x)\psi_1(y)u_x + \phi_2(x)\psi_2(y)u_y = f, \quad (x, y) \in \Omega, \quad (3.5)$$

with $\varepsilon > 0$, and for the sake of simplicity, $\Omega = [0, 1] \times [0, 1]$ with zero Dirichlet boundary conditions. Using standard centered finite difference discretization for each term, and setting $\mathbf{U}_{ij} := u(x_i, y_j)$, where (x_i, y_j) are interior grid nodes, $i, j = 1, \dots, n$, we obtain

$$T\mathbf{U} + \mathbf{U}T + \Phi_1 B \mathbf{U} \Psi_1^\top + \Psi_2 \mathbf{U} (\Phi_2 B)^\top = F, \quad F = (f(x_i, y_j)); \quad (3.6)$$

here

$$T = -\frac{\varepsilon}{h^2} \text{tridiag}(1, -2, 1), \quad B = \frac{1}{2h} \text{tridiag}(-1, 0, 1),$$

and

$$\Phi_k = \text{diag}(\phi_k(x_1), \dots, \phi_k(x_n)), \quad \Psi_k = \text{diag}(\psi_k(y_1), \dots, \psi_k(y_n)), \quad k = 1, 2,$$

where h is the mesh size. Equation (3.6) is a four-term linear matrix equation in \mathbf{U} and it was used in the early literature on difference equations; we refer the reader to, e.g., [41], for similar derivations. Common strategies then transform the problem above into the following standard non-symmetric linear system by means of the Kronecker product:

$$\left(I \otimes T + T \otimes I + \Psi_1 \otimes (\Phi_1 B) + (\Phi_2 B) \otimes \Psi_2 \right) \mathbf{u} = \tilde{f}, \quad \mathbf{u} := \text{vec}(U), \tilde{f} = \text{vec}(F),$$

for whose solution a vast literature is available. We are unaware of recent strategies that exploit the matrix equation formulation of the problem for its numerical solution, whereas the matrix structure may suggest particular preconditioning strategies.

In the context of dynamical system analysis, multi-term matrix equations of the type (1.2) arise in the numerical treatment of bilinear systems in the form (see, e.g., [116],[215])

$$\dot{x}(t) = (A + u(t)N)x(t) + Bu(t), \quad x(0) = x_0, \quad y(t) = Cx(t), \quad (3.7)$$

which occur when the model accounts for a stochastic component by means of the term involving N . Other generalizations of Gramians can thus be considered, which can be written as the solution \mathbf{X} to the following multi-term linear matrix equation,

$$A\mathbf{X} + \mathbf{X}A^\top + N\mathbf{X}N^\top + BB^\top = 0,$$

together with its counterpart with respect to $C^\top C$; note that extra terms N_i can be included in the sum; see [30] and references therein. The solution \mathbf{X} carries information on the reachability and observability properties of the state vectors [105]. The one above is an example of linear jump systems (see [181]), in which the linear coefficient matrices depend on a Markovian random process, giving rise to systems of matrix equations with the extra term, accounting for the probabilistic nature of the problem.

Another typical emerging setting where the multi-term matrix equation in (1.2) arises is the analysis of uncertainty quantification in data modelling. For instance, the stochastic steady state diffusion equation with homogeneous Dirichlet boundary conditions is given by

$$\begin{cases} -\nabla \cdot (c\nabla p) = f & \text{in } D \times \Omega \\ p = 0 & \text{on } \partial D \times \Omega, \end{cases} \quad (3.8)$$

where D is a sufficiently regular spatial domain and Ω is a probability sample space. Both the forcing term f and the diffusion coefficient c have a stochastic component. By properly discretizing the *weak* formulation of (3.8), and under certain assumptions on the stochastic discretized space, one obtains the algebraic linear system (see, e.g., [83] and references therein):

$$\mathcal{A}\mathbf{p} = f, \quad \mathcal{A} = G_0 \otimes K_0 + \sum_{r=1}^m \sqrt{\lambda_r} G_r \otimes K_r. \quad (3.9)$$

By passing to the matrix formulation, and introducing the matrix \mathbf{X} of coefficients in \mathbf{p} , (3.9) can be rewritten as

$$K_0 \mathbf{X} G_0^\top + \sum_{r=1}^m \sqrt{\lambda_r} K_r \mathbf{X} G_r^\top = F, \quad (3.10)$$

where F contains the components of f , in which each column of F corresponds to a different basis element in the probability space. In many simulations, while the underlying mathematical formulation is still (3.8), the quantity of interest is $c\nabla p$, rather than p . Using for instance the derivation in [87], a direct approximation to $c\nabla p$ is obtained by introducing the variable (flux) $\vec{u} = c\nabla p$ which gives

$$\begin{cases} c^{-1}\vec{u} - \nabla p = 0 & \text{in } D \times \Omega \\ -\nabla \cdot \vec{u} = f & \text{in } D \times \Omega, \end{cases} \quad p = 0 \quad \text{on} \quad \partial D \times \Omega. \quad (3.11)$$

By means of a discretization with proper (tensor products of) finite element spaces of the weak formulation of (3.11) (see, e.g., [87],[83],[201]), one obtains the following saddle point algebraic linear system

$$\begin{bmatrix} A & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix} = \begin{bmatrix} 0 \\ f \end{bmatrix}, \quad A = G_0 \otimes K_0 + \sum_{r=1}^m \sqrt{\lambda_r} G_r \otimes K_r, \quad B = G_0 \otimes B_0. \quad (3.12)$$

The solution vectors \mathbf{u} and \mathbf{p} contain the two dimensional coefficients of the (discrete) expansions of \vec{u} , and p column by column. Once again, a closer look at the two equations above reveals that the matrix formulation could replace the Kronecker products. Indeed, if \mathbf{U} is the matrix such that $\mathbf{u} = \text{vec}(\mathbf{U})$, whose coefficients are $(u_{j\ell})$, and similarly for \mathbf{P} , then the linear system above reads:

$$K_0 \mathbf{U} G_0^T + \sum_{r=1}^m \sqrt{\lambda_r} K_r \mathbf{U} G_r^T + B_0^T \mathbf{P} G_0 = 0, \quad (3.13)$$

$$B_0 \mathbf{U} G_0^T = F, \quad (3.14)$$

with obvious meaning for F . This system is a natural generalization of the case in (3.10), and may be thought of as a saddle point *generalized matrix* system. This type of systems of linear matrix equations will be discussed in section 7.2.

4. Continuous-time Sylvester equation. The continuous-time Sylvester equation is possibly the most broadly employed linear matrix equation, and is given as

$$A\mathbf{X} + \mathbf{X}B = C, \quad (4.1)$$

with $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{m \times m}$ and $C \in \mathbb{R}^{n \times m}$. In general, the dimensions of A and B may be orders of magnitude different, and this fact is key in selecting the most appropriate numerical solution strategy.

A general result on the consistency of the Sylvester equation was given by Roth in 1952 in [209], and it reads: *the equation (4.1) admits a solution if and only if the matrices*

$$\begin{bmatrix} A & -C \\ 0 & -B \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} A & 0 \\ 0 & -B \end{bmatrix} \quad (4.2)$$

are similar; the similarity transformation matrix is given by

$$\begin{bmatrix} I & \mathbf{X} \\ & I \end{bmatrix}$$

where \mathbf{X} is the solution to (4.1).

Using the Kronecker product, the matrix equation in (4.1) can be rewritten as the following standard (vector) linear system

$$\mathcal{A}\mathbf{x} = c, \quad \text{with} \quad \begin{aligned} \mathcal{A} &= I_m \otimes A + B^\top \otimes I_n \\ \mathbf{x} &= \text{vec}(\mathbf{X}), \quad c = \text{vec}(C), \end{aligned} \quad (4.3)$$

from which we can deduce that the system admits a solution for any c and this is unique, if and only if the matrix \mathcal{A} is nonsingular. Taking into account Lemma 2.1(iii), this is equivalent to requiring that $\text{spec}(A) \cap \text{spec}(-B) = \emptyset$ (see, e.g., [130, Th. 4.4.6]). In the following we shall thus always assume that this latter condition is satisfied, so that the solution to (4.1) exists and is unique; standard matrix analysis books describe the case when this spectral condition is not satisfied (see, e.g., [130], [165]). The homogeneous case, namely when $C = 0$, can be handled correspondingly: the matrix equation has only the trivial solution $\mathbf{X} = 0$ if and only if $\text{spec}(A) \cap \text{spec}(-B) = \emptyset$ [95, sec.17.8].

The solution \mathbf{X} of (4.1) may be written in closed form in a number of different ways. These forms have been derived in different references throughout the 1950's and 1960's, with contributions by E. Heinz, A. Jameson, M.G. Krein, E.C. Ma, M. Rosenblum, W. E. Roth, etc. A beautiful account of these early contributions can be found in the survey by P. Lancaster [166], to which we refer the reader also for the bibliographic references. Here we report the main closed forms:

- (a) *Integral of resolvents.* The following representation, due to Krein, exploits spectral theory arguments:

$$\mathbf{X} = -\frac{1}{4\pi^2} \int_{\Gamma_1} \int_{\Gamma_2} \frac{(\lambda I_n - A)^{-1} C (\mu I_m - B)^{-1}}{\lambda + \mu} d\mu d\lambda, \quad (4.4)$$

where Γ_1, Γ_2 are contours containing and sufficiently close to, the spectra of A and B , respectively.

- (b) *Integral of exponentials.* This representation, due to Heinz, is tightly connected to the previous one,

$$\mathbf{X} = -\int_0^\infty e^{At} C e^{Bt} dt, \quad (4.5)$$

where e^{Ht} is the matrix exponential of Ht . Here the spectra of A and B are supposed to be separated by a vertical line.

- (c) *Finite power sum.* Let $C = C_A C_B^\top$. Let a_m of degree m be the minimal polynomial of A with respect to C_A , namely the smallest degree monic polynomial such that $a_m(A)C_A = 0$. Analogously, let b_k of degree k be the minimal polynomial of B with respect to C_B . Then

$$\begin{aligned} \mathbf{X} &= \sum_{i=0}^{m-1} \sum_{j=0}^{k-1} \gamma_{ij} A^i C B^j \\ &= [C_A, AC_A, \dots, A^{m-1}C_A](\gamma \otimes I) \begin{bmatrix} C_B^\top \\ C_B^\top B \\ \vdots \\ C_B^\top B^{k-1} \end{bmatrix} \end{aligned} \quad (4.6)$$

where γ is the solution of the Sylvester equation with coefficient matrices the companion matrices of a_m and b_k , and right-hand side the matrix

$[1; 0 \dots; 0][1, 0, \dots, 0]$ [67]; a block version of this result using minimal *matrix* polynomials can also be derived [223].

- (d) *Similarity transformations.* Strictly related to (c), in addition this form assumes that A and B can be diagonalized, $U^{-1}AU = \text{diag}(\lambda_1, \dots, \lambda_n)$ and $V^{-1}BV = \text{diag}(\mu_1, \dots, \mu_m)$. Let $\tilde{C} = U^{-1}CV$. Then

$$\mathbf{X} = U\tilde{X}V^{-1}, \quad \text{with} \quad \tilde{x}_{ij} = \frac{\tilde{c}_{ij}}{\lambda_i + \mu_j}.$$

Other representations can be found in [166], and for more general equations of the same type, in [261]. We also mention that the columns of $[\mathbf{X}; I]$ span an invariant subspace for the left matrix in (4.2), that is

$$\begin{bmatrix} A & -C \\ 0 & -B \end{bmatrix} \begin{bmatrix} \mathbf{X} \\ I \end{bmatrix} = \begin{bmatrix} \mathbf{X} \\ I \end{bmatrix} S \quad (4.7)$$

where it holds that $S = -B$. Equation (4.7) has been used both to derive matrix properties of the solution \mathbf{X} , and also to construct solution devices.

In [67] the closed form in (c) is used to derive results on the solution rank; Results on the nonsingularity of the solution based on the same conditions are also contained in [118]. For more general equations, corresponding nonsingularity conditions can be found, e.g., in [260]. In [67], the controllability (resp. observability) of the pair (A, C_A) (resp. (B^T, C_B)) plays a crucial role³.

Early computational methods relied on one of the analytic expressions above; see the account on early computational methods in [88]. However, these closed forms are no longer used to numerically solve the Sylvester equation, as they are clearly inefficient - and possibly unstable - even for small n, m . On the other hand, they have been used as motivation to several successful methods and they represent an important starting point for theoretical investigations of numerical approaches.

4.1. Stability and sensitivity issues of the Sylvester equation. In this section we provide a brief account of the sensitivity issues encountered when solving the Sylvester equation. The topic is broad, and it also involves the solution of related matrix equations; we refer to the thorough treatment in [155] for a full account on the perturbation theory of this and other important equations in control.

The sensitivity to perturbations of the solution \mathbf{X} to (4.1) is inversely proportional to the separation between A and $-B$, where the separation function of two matrices A_1 and A_2 is defined as

$$\text{sep}_p(A_1, A_2) = \min_{\|P\|_p=1} \|A_1P - PA_2\|_p,$$

with $p = 2, F$; see, e.g., [236]. This can be viewed by recalling that the columns of $[\mathbf{X}; I]$ are a basis for an invariant subspace for the first block matrix in (4.7). We refer the reader to, e.g., [96, section 7.6.3] where the role of $\|\mathbf{X}\|$ in the conditioning of the associated eigenvalues is emphasized. More specifically, it holds that

$$\|\mathbf{X}\| < 2 \frac{\|C\|}{\text{sep}_2(A, -B)}. \quad (4.8)$$

³A pair (M, C) is controllable if the matrix $[C, MC, \dots, M^{n-1}C]$ has full row rank n , equal to the row dimension of M ; (M, C^T) is observable if (M^T, C) is controllable.

For non-normal matrices, the bound above suggests that a good spectral distance between A and $-B$ might not be sufficient to limit the size of $\|\mathbf{X}\|$, since $\text{sep}_2(A, -B)$ can be much smaller than the distance between the spectra of A and $-B$. The function sep plays the role of a condition number for the following Sylvester operator

$$S : \mathbb{R}^{n \times m} \rightarrow \mathbb{R}^{n \times m}, \quad S(X) = AX + XB; \quad (4.9)$$

numerical estimates for the sep function can be obtained by carefully adapting classical strategies [49]. The occurrence of the sep function in the bound (4.8) suggests that for small scale equations, algorithms that rely on orthogonal reduction should be preferred in terms of numerical stability. Methods that rely on more general transformations $\tilde{\mathbf{X}} = U\mathbf{X}V^{-1}$ may transfer the ill-conditioning of the transformation matrices U and V onto large errors in the obtained solution; moreover, (see, e.g., [236, Exercise V.2.1])

$$\frac{\text{sep}(A, B)}{\kappa(U)\kappa(V)} \leq \text{sep}(UAU^{-1}, VB^{-1}V) \leq \kappa(U)\kappa(V)\text{sep}(A, B).$$

A major difference between matrix equations and standard linear systems lies in their stability properties. In particular, a small Sylvester equation residual does not necessarily imply a small backward error [122, sec.15.2]. Define the backward error for an approximation \mathbf{X} as

$$\eta(\mathbf{X}) := \min\{\varepsilon : (A + \Delta A)\mathbf{X} + \mathbf{X}(B + \Delta B) = C + \Delta C, \\ \|\Delta A\|_F \leq \varepsilon\|A\|_F, \|\Delta B\|_F \leq \varepsilon\|B\|_F, \|\Delta C\|_F \leq \varepsilon\|C\|_F\},$$

and the residual as $R = C - (A\mathbf{X} + \mathbf{X}B)$. Then ([122])

$$\eta(\mathbf{X}) \leq \mu \frac{\|R\|_F}{(\|A\|_F + \|B\|_F)\|\mathbf{X}\|_F + \|C\|_F}, \quad (4.10)$$

where μ is an amplification factor depending on the data norms and on the singular values of \mathbf{X} . For instance, for $n = m$ this factor has the expression

$$\mu = \frac{(\|A\|_F + \|B\|_F)\|\mathbf{X}\|_F + \|C\|_F}{((\|A\|_F^2 + \|B\|_F^2)\sigma_{\min}(\mathbf{X})^2 + \|C\|_F^2)^{\frac{1}{2}}},$$

making the dependence on the norm and ill-conditioning of \mathbf{X} more apparent. A more complex situation occurs for $n \neq m$; we refer the reader to [123, sec.15.2] for more details, and to [155] for a more general treatment. We also mention that in [244] bounds for the norm of the solution \mathbf{X} and of its perturbation are obtained, that emphasize the influence of the possibly low-rank right-hand side on the sensitivity of the solution itself. The distribution of the singular values of \mathbf{X} plays a crucial role in the stability analysis of dynamical systems and also in the quality of low rank approximations. In section 4.4 we recall same available estimates for the singular values, that also motivate the development of low rank approximation methods.

4.2. Sylvester equation. Small scale computation. A robust and efficient method for numerically solving Sylvester equations of small and moderate size was introduced in 1972 by Bartels and Stewart [15], and with some modifications is still the state-of-the-art; in section 8 we give an account of current software, highly relying on this method. The idea is to compute the Schur decomposition of the two coefficient matrices and then transform the given equation into an equivalent one that uses

the quasi-lower/upper structure of the Schur matrices. The last equation can then be explicitly solved element by element. For introducing the algorithm, let us first consider the general case of complex A and B . Then the following steps are performed (see, e.g., [96]):

Algorithm 2.

1. Compute the Schur forms: $A^* = URU^*$, $B = VSV^*$ with R, S upper triangular;
2. Solve $R^*\mathbf{Y} + \mathbf{Y}S = U^*CV$ for \mathbf{Y} ;
3. Compute $\mathbf{X} = U\mathbf{Y}V^*$.

The Schur forms in the first step are obtained by the QR iteration [96], while the third step is a simple product. It remains to explain how to solve the new structured Sylvester equation in the second step. Since R^* is lower triangular and S is upper triangular, the (1,1) element of \mathbf{Y} can be readily obtained. From there the next elements of the first row in \mathbf{Y} can also be obtained sequentially. A similar reasoning can be used for the subsequent rows.

In the case of real A and B , the *real* Schur form may be exploited, where R and S are now quasi-triangular, that is the diagonals have 2×2 and 1×1 blocks, corresponding to complex and real eigenvalues, respectively. The process relies on the equivalence between a 2×2 Sylvester equation and the associated Kronecker form in (4.3) [221]. The same sequential process as in the complex case can be employed to compute the elements of \mathbf{Y} , as long as the diagonal blocks can be made conforming [15],[96]. The method just outlined is at the core of most linear matrix equations solvers in software packages such as LAPACK⁴ and SLICOT [247],[230],[33]. The leading computational cost is given by the Schur forms in step one, which for real matrices are nowadays performed in real arithmetic. Explicitly writing the Schur form costs at least $10n^3$ for a matrix of size n [96]; to limit costs, the Bartels-Stewart algorithm is commonly employed only if either A or B is already in Schur or upper Hessenberg form; see, e.g., [230]. For general matrices A and B , the method proposed by Golub, Nash and Van Loan in 1979 ([97]) can be considerably faster, especially if either m or n is significantly smaller than the other. This latter method replaces the Schur decomposition of the larger matrix, say, B , with the Hessenberg decomposition of the same matrix whose computational cost is $5/3m^3$, which should be compared with $10m^3$ of the Schur form [97]. We refer the reader to [221, sec.2.3.1] for a more detailed comparison on the computational costs. In [233], a variant of the Bartels-Stewart algorithm is proposed: the forward-backward substitution in step 2 is performed by a column-wise block scheme, which seems to be better suited for modern computer architectures than the original complex version. In [141],[142], the authors propose an even more effective implementation, based on splitting the matrices - already in block triangular form, and then recursively solving for each block. For instance, if A is much larger than B ($n \geq 2m$), then the original equation can be written as

$$\begin{bmatrix} A_{11} & A_{12} \\ & A_{22} \end{bmatrix} \begin{bmatrix} X_1 \\ X_2 \end{bmatrix} + \begin{bmatrix} X_1 \\ X_2 \end{bmatrix} B = \begin{bmatrix} C_1 \\ C_2 \end{bmatrix},$$

with obvious meaning for the blocks. The second block equation gives the smaller size Sylvester equation $A_{22}X_2 + X_2B = C_2$, which can again be split by using the block triangular form of A_{22} , and the solution obtained in a recursive manner. Once

⁴<http://www.netlib.org/lapack/>

X_2 is fully recovered, X_1 can be computed by recursively solving with the updated right-hand side in the first block equation above. Different size cases and different triangular structures can be handled and are described in [141]. These advanced strategies have been included in the software package RECSY⁵ and in LAPACK; see section 8.

Iterative solution strategies for small size matrices have also been proposed: given an initial guess \mathbf{X}_0 , they determine a sequence of matrices $\mathbf{X}_1, \dots, \mathbf{X}_k, \dots$ that converge to \mathbf{X} . These are related to a basic Newton iteration for approximating the matrix sign function. In section 5.2.3 we will give more details in relation with the Lyapunov equation, although the procedure can be used for stable Sylvester equations as well [38]. These approaches are easier to parallelize than QR based methods. For instance, it is shown in [38] that they provide high efficiency and scalability on clusters of processors.

To conclude, a special mention should be paid to the Sylvester equation with $B = -A$, yielding the so-called *displacement equation*

$$A\mathbf{X} - \mathbf{X}A = C, \quad (4.11)$$

which measures how far A and \mathbf{X} are from commuting; see, e.g., [94] for typical applications in the context of structured matrices such as Cauchy-like and Toeplitz matrices.

4.3. Sylvester equation. Large A and small B . When either n or m is large, Schur factorization may require a prohibitive amount of space, due to the dense nature of the corresponding large matrix. Selecting the most appropriate solver still depends on whether the smaller matrix has very small dimension. Different approaches can then be used when decomposing the small matrix is feasible⁶. To fix ideas, and without loss of generality, we shall assume that B is small (size less than 1000) and A is large (size much bigger than 1000), so that $m \ll n$.

In this section we thus consider that the equation can be visualized as:

$$\begin{bmatrix} A \end{bmatrix} \begin{bmatrix} \mathbf{X} \end{bmatrix} + \begin{bmatrix} \mathbf{X} \end{bmatrix} \begin{bmatrix} B \end{bmatrix} = \begin{bmatrix} C \end{bmatrix}, \quad (4.12)$$

so that the large dimension of A makes the methods discussed in section 4.2 unfeasible. This setting arises for instance in the solution of eigenvalue problems [259, sec.2.4, sec.6.6] and in (separable) boundary value problems [256],[257],[41]. We readily notice that for very small m , the transformation with the Kronecker product (4.3) may be appealing, since the dimension of the linear system may be just a few (m) times that of A . However, projection methods acting on the original matrix equation turn out to be extremely effective in this case, possibly explaining the fewer attempts to pursue such Kronecker formulation. We next describe some of the standard approaches currently employed in the literature and in applications.

Assume that B can be spectrally decomposed cheaply and stably. Then by writing $B = WSW^{-1}$ with $S = \text{diag}(s_1, \dots, s_m)$, we obtain

$$A\widehat{\mathbf{X}} + \widehat{\mathbf{X}}S = \widehat{C}, \quad \widehat{\mathbf{X}} = \mathbf{X}W, \quad \widehat{C} = CW. \quad (4.13)$$

⁵<http://www8.cs.umu.se/~isak/recsy/>.

⁶Feasibility is machine architecture dependent; nonetheless, a matrix of dimension much less than one thousand should be considered small.

For B Hermitian, $W^{-1} = W^*$. Each column of $\widehat{\mathbf{X}}$ can be obtained by solving a shifted linear system $(A + s_i I)(\widehat{\mathbf{X}})_i = (\widehat{C})_i$, where $(\widehat{\mathbf{X}})_i$ denotes the i th column of $\widehat{\mathbf{X}}$. The main steps can be summarized in the following algorithm:

Algorithm 3.

1. Compute the decomposition $B = WSW^{-1}$
2. Set $\widehat{C} = CW$
3. For each i , solve $(A + s_i I)(\widehat{\mathbf{X}})_i = (\widehat{C})_i$
4. Compute $\mathbf{X} = \widehat{\mathbf{X}}W^{-1}$

The shifted systems in step 3 can be solved simultaneously by using standard solvers for algebraic linear systems, either direct or iterative; see, e.g., [212], [228] and their references. We also note that step 3 is “embarrassingly parallel” in case different systems can be distributed on a multiprocessor machine.

If the eigendecomposition of B is not appealing, then one can resort to a (complex) Schur decomposition $B = QR_BQ^*$, giving $A\mathbf{X}Q + \mathbf{X}QR_B = CQ$. Since R_B is upper triangular, these systems can still be solved by using the shifted form, but this time in sequence: letting r_{ij} be the (i, j) entry of R_B and $\widehat{C} = CQ$, we have

$$\text{for } i = 1, \dots, m, \quad (A + r_{ii}I)(\widehat{\mathbf{X}})_i = (\widehat{C})_i - \sum_{k=1}^{i-1} r_{ki}(\widehat{\mathbf{X}})_k, \quad \widehat{\mathbf{X}} = \mathbf{X}Q; \quad (4.14)$$

such an approach has been used in different contexts, see, e.g., [108],[234],[32], where the considered Sylvester equation is occasionally called a *sparse-dense* equation.

For moderate n , the use of direct methods in (4.13) and (4.14) may entail the use of complex arithmetic if the shifts (eigenvalues) are complex, significantly increasing the computational cost; the alternative of solving two real systems also leads to higher computational costs. In addition, in case the use of sparse direct methods appears to be competitive, it should be noticed that only the sparsity analysis step can be done once for all, whereas the actual decomposition needs to be performed again for each distinct shift.

Major computational savings may be obtained if C is low rank, namely $C = C_0R$, with $C_0 \in \mathbb{R}^{n \times \bar{m}}$ and $\bar{m} < m$. Indeed, the m shifted systems can be solved more efficiently by only working with the common matrix C_0 . For the rest of this section we assume that C is full rank, and postpone the treatment of the low-rank case to later, when we discuss the occurrence of large B . Indeed, the rank of C is key in developing general projection methods, as explained next.

Projection methods. Let \mathcal{V} be a subspace⁷ of \mathbb{C}^n of dimension k , and let the columns of $V_k \in \mathbb{C}^{n \times k}$ span \mathcal{V} . An approximate solution \mathbf{X}_k with $\text{range}(\mathbf{X}_k) \subset \mathcal{V}$ is sought such that

$$R_k := A\mathbf{X}_k + \mathbf{X}_k B - C \approx 0.$$

Several options arise, depending on the choice of \mathcal{V} and on the strategy to determine \mathbf{X}_k within the space \mathcal{V} . For a given \mathcal{V} , let thus $\mathbf{X}_k = V_k \mathbf{Y}_k \approx \mathbf{X}$, for some $\mathbf{Y}_k \in$

⁷We use complex arithmetic for \mathcal{V} to allow for complex spaces also for real data, which may occur when using Rational Krylov subspaces with complex shifts. A careful implementation can construct a real space in case conjugate shifts are used. For the sake of generality we stick to complex arithmetic for \mathcal{V} .

$\mathbb{R}^{k \times m}$ to be determined. Recalling the operator \mathcal{S} defined in (4.9), we observe that \mathcal{S} generalizes to the “block” B the concept of shifted matrices, namely

$$x \mapsto (A + \beta I)x = Ax + x\beta.$$

Therefore, it is very natural to extend the algorithmic strategies of linear systems to the case of \mathcal{S} . Extensions of the linear system solvers CG (FOM) and MINRES (GMRES) can be thought of for A Hermitian (nonhermitian), although the actual implementation differs. All these solvers are derived by imposing some orthogonality condition on the system residual. If we require that the columns of the matrix R_k be orthogonal to the approximation space \mathcal{V} in the Euclidean inner product, then we are imposing the following Galerkin condition (see also (4.23)):

$$V_k^* R_k = 0 \quad \Leftrightarrow \quad (I \otimes V_k)^* \text{vec}(R_k) = 0.$$

For simplicity let us assume that $V_k^* V_k = I$. Then

$$0 = V_k^* R_k = V_k^* A V_k \mathbf{Y}_k + \mathbf{Y}_k B - V_k^* C. \quad (4.15)$$

The condition thus gives a new, reduced in size, Sylvester equation to be solved. Under the hypothesis that $\text{spec}(V_k^* A V_k) \cap \text{spec}(-B) = \emptyset$, equation (4.15) can be solved efficiently by one of the methods discussed in section 4.2. The procedure above holds for any space \mathcal{V} and associated full rank matrix V_k . Therefore, the effectiveness of the approximation process depends on the actual selection of \mathcal{V} . A well exercised choice is given by the block Krylov subspace

$$K_k^\square(A, C) = \text{range}([C, AC, \dots, A^{k-1}C]). \quad (4.16)$$

The following result proved in [207, Lemma 2.1], [223] generalizes the well known shift invariance property of vector Krylov subspaces to the case of blocks, where the $m \times m$ matrix B plays the role of the shift.

PROPOSITION 4.1. *Define $\mathcal{S}^j(C) = \mathcal{S}(\mathcal{S}^{j-1}(C))$, $j > 0$ and $\mathcal{S}^0(C) = C$. Then*

$$K_k^\square(A, C) = K_k^\square(\mathcal{S}, C) := \text{range}([C, \mathcal{S}(C), \dots, \mathcal{S}^{k-1}(C)]).$$

With the space in (4.16), the procedure just outlined is the complete analog of the one giving rise to the Full Orthogonalization Method (FOM) for $m = 1$ or for $B = 0$. However, due to possible loss of rank in the basis, it was suggested in [207] to generate the subspace with A rather than with \mathcal{S} . As an example, Algorithm 4 describes an implementation of the projection method with the generation of the block Krylov subspace and the determination of the approximation by imposing the Galerkin orthogonality condition.

Algorithm 4. Given A, B, C

1. Orthogonalize the columns of C to get $v_1 = V_1$
2. $k = 1, 2, \dots$
3. Compute \mathbf{Y}_k , solution to $(V_k^* A V_k) \mathbf{Y} + \mathbf{Y} B - V_k^* C = 0$
4. If converged $\mathbf{X}_k = V_k \mathbf{Y}_k$ and stop
5. Arnoldi procedure for the next basis block:
 - $\hat{v} = A v_k$
 - Make \hat{v} orthogonal wrto $\{v_1, \dots, v_k\}$
 - Orthogonalize (wrto 2-norm) the columns of \hat{v} to get v_{k+1}
 - Update: $V_{k+1} = [V_k, v_{k+1}]$

For later reference, we remark that the Arnoldi procedure used in Algorithm 4 generates a matrix recurrence that can be written as

$$AV_k = V_k H_k + \hat{v} e_k^\top, \quad (4.17)$$

where \hat{v} is the new block of basis vectors, prior orthogonalization, and H_k contains the orthogonality coefficients, with $H_k = V_k^* AV_k$.

One could consider constraint spaces different from the approximation spaces; in this case, a so-called Petrov-Galerkin condition is imposed on the residual. To this end, let us consider the matrix inner product defined as

$$\langle Y, X \rangle_F = \text{trace}(Y^* X), \quad X, Y \in \mathbb{R}^{n \times m}. \quad (4.18)$$

Following the standard linear system case with $m = 1$, and using, e.g., the space spanned by the columns of AV_k , one would be tempted to impose the condition $(AV_k)^* R_k = 0$ in the Euclidean inner product, giving

$$V_k^* A^* AV_k \mathbf{Y}_k + V_k^* A^* V_k \mathbf{Y}_k B - V_k^* A^* C = 0. \quad (4.19)$$

In the standard ($B = 0$) linear system setting, this condition is equivalent to minimizing the residual R_k in the Frobenius norm, that is

$$\min_{\mathbf{Y}_k \in \mathbb{R}^{k \times m}} \|R_k\|_F. \quad (4.20)$$

However, for $B \neq 0$, such equivalence does *not* hold, that is, the solution to (4.19) is not a residual minimizing approximation. To attain a residual minimization, the orthogonality condition should be applied to the operator \mathcal{S} in (4.9) in the Frobenius inner product (4.18); to this end, we note that the adjoint operator \mathcal{S}^* with respect to the inner product in (4.18) is given by $\mathcal{S}^*(X) = A^* X + X B^*$.

PROPOSITION 4.2. [207, sec.3] *Let $\mathbf{Y}_k \in \mathbb{R}^{k \times m}$, and let $R_k = AV_k \mathbf{Y}_k + V_k \mathbf{Y}_k B - C$ be the associated residual. Then*

$$\mathbf{Y}_k = \arg \min_{\mathbf{Y}_k \in \mathbb{R}^{k \times m}} \|R_k\|_F \quad \text{if and only if} \quad R_k \perp_F \mathcal{S}(K_m(\mathcal{S}, V_1)).$$

For the choice $\mathcal{V} = K_k^\square(A, C)$, the minimization process in (4.20) is the matrix analog of GMRES (for $m = 1$ or $B = 0$) (see [212, sec. 6.12]). Similar results are discussed independently in [108]. Inspired by the ‘‘block shift’’ invariance of Proposition 4.1, the authors of [207] provide a detailed description of the parallel between solving (4.12) for $m \ll n$ with Galerkin and with minimizing procedures, and solving linear systems $A\mathbf{X} = C$ by means of block methods. Upper bounds for the residual norm of Galerkin and residual minimizing methods with $\mathcal{V} = K_k^\square(A, C)$ are also provided in [207], together with numerical experiments on the performance of the approaches.

Preconditioned global Krylov subspaces have also been proposed as approximation space [46], which however simply amount to a convenient implementation of a subspace method for the Kronecker formulation of the problem; see also section 4.4.1.

An alternative choice of approximation space \mathcal{V} has recently shown great potential compared with the block Krylov subspace, and it is given by the *Extended* Krylov subspace, defined as

$$\mathbf{EK}_k(A, C) := K_k^\square(A, C) + K_k^\square(A^{-1}, A^{-1}C). \quad (4.21)$$

Since the spaces are nested, namely $\mathbf{EK}_k(A, C) \subseteq \mathbf{EK}_{k+1}(A, C)$, the space can be generated iteratively, allowing one to improve the approximate solution as the recurrence proceeds. For A large and sparse, experiments in [225] show that the good performance of the derived method seems to fully compensate the high costs of solving linear systems with A at each iteration.

4.4. Sylvester equation. Large A and large B . In the most general case, both A and B have large dimensions. This setting arises in many situations, as in the discretization of separable PDEs [78], or in the computation of the *cross* Gramian in control [4]. A particularly important observation is that the dimensions of A and B determine that of \mathbf{X} , and that although A and B may be sparse, \mathbf{X} is dense, in general. In this context, the distribution of the singular values of \mathbf{X} plays a key role in the development and convergence analysis of iterative solution methods. Indeed, a Sylvester equation having solution with exponentially decaying singular values can be well approximated by a low rank matrix. The possibility of writing $C = C_1 C_2^\top$ with C_1, C_2 with low column rank is crucial to obtain good low-rank approximations to \mathbf{X} , thus avoiding the storage of the whole matrix, which is in general prohibitive. We recall here the result described by Sabino in [213, Th.2.1.1], while Sabino's PhD thesis contains further discussion related to this bound. Here K and K' are the complete elliptic integrals of the first kind⁸ [1]. Additional considerations and results are postponed to the Lyapunov equation case in section 5.

THEOREM 4.3. *Let A and B be stable and real symmetric, with spectra contained in $[a, b]$ and $[c, d]$, respectively. Define $\eta = 2(b - a)(d - c)/((a + c)(b + d))$. Assume C is of rank p . Then the singular values $\sigma_1 \geq \dots \geq \sigma_{\min\{m, n\}}$ of the solution \mathbf{X} to (4.1) satisfy*

$$\frac{\sigma_{pr+1}}{\sigma_1} \leq \left(\frac{1 - \sqrt{k'_r}}{1 + \sqrt{k'_r}} \right)^2, \quad 1 \leq pr < n,$$

where $k'_r = 1/(1 + \eta + \sqrt{\eta(\eta + 2)})$ is the complementary elliptic modulus corresponding to the nome q^r , $q := \exp(-\pi K'/K)$.

A more accessible and practical estimate for $B = A$ and small condition number $\kappa(A) = \|A\| \|A^{-1}\|$ may be obtained as ([213])

$$\frac{\sigma_{pr+1}}{\sigma_1} \lesssim 4 \exp(-\pi^2 r / \log(4\kappa(A))). \quad (4.22)$$

Easy to use variants of (4.22) in [213] are favorably compared with earlier estimates in [197]. Results for A and B nonsymmetric are few; non-normality may strongly influence the solution rank, so that results will significantly depart from the above bound. A satisfactory understanding of the singular value decay for non-normal coefficient matrices is still lacking.

From a numerical analysis viewpoint, we notice that the main rational approximation ingredients used for results of the type above are the same as those obtained for rational space projections and ADI-type iterations (see section 4.4.1 and section 4.4.2, respectively), which also rely on minimax rational approximations; in fact, the result above is intimately related to similar estimates by Ellner and Wachspress for ADI in [78], [79].

⁸They are defined as $K = K(k) = \int_0^1 [(1 - t^2)(1 - kt^2)]^{-1/2} dt$ and $K' = K(1 - k)$, with k being the modulus, $k = \sqrt{1 - (k')^2}$, while the complementary elliptic modulus k' is given.

Numerical methods in the literature have mainly proceeded in three directions: Projection type approaches (mostly based on the Krylov subspace family), matrix updating sequences (such as Alternating-Direction-Implicit iterations), and sparse data format recurrences. Combinations of these have also been explored.

The convergence rates of the strategies in the first two classes strongly depend on the spectral properties of the coefficient matrices (eigenvalues or field of values). For those problems with unfavorable spectral information, for instance a large field of values close to the origin, the most efficient available methods rely on iterations that involve solving linear systems at each step, either with A or with $A + \sigma I$ for some appropriately chosen σ . For A large but very sparse, these solves can be conveniently carried out by means of direct methods. On the other hand, if the direct solution with A becomes prohibitively expensive, in terms of CPU time or memory requirements, an (inner) iterative solution of the linear systems with $A + \sigma I$ is performed at each step, giving rise to an inner-outer procedure. In this case, one usually talks about “inexact solves”, unless the iterative process allows one to reach machine precision accuracy. Such a consideration noticeably influences the evaluation of the computational costs of these methods, whose performance is thus problem dependent.

Due to the important role the Lyapunov equation has in control problems, many authors have developed numerical procedures specifically for the Lyapunov equation, and not for the Sylvester equation, although in many cases they could be extended to the latter in a natural manner. For historical reasons, and also to avoid constant reference to the equation context, we will refer to the literature the way the methods were originally presented. In particular, it will be apparent that the literature on large scale Lyapunov equations is richer than that for the Sylvester equation, especially in the large scale case.

We also notice that, as a major distinction from linear vector equations, the numerical solution of matrix equations cannot directly rely on preconditioning strategies, unless the Kronecker formulation is employed. Indeed, preconditioning methods would necessarily destroy the symmetry properties of the problem, which allows one to deal with computational costs that depend on powers of n and m , but not on powers of $n \cdot m$. As an example, let us assume that a nonsingular matrix P exists⁹ such that $P^{-1}A$ and $P^{-1}B^T$ have better spectral properties than the original matrices - for A, B symmetric, this requirement corresponds to a better clustering of the eigenvalues. Then we could consider applying P as follows

$$P^{-1}A\mathbf{X}P^{-T} + P^{-1}\mathbf{X}BP^{-T} = P^{-1}CP^{-T}.$$

To be able to rewrite such an equation in terms of a single unknown matrix, one could premultiply and postmultiply \mathbf{X} by P^{-1} and P^{-T} , that is

$$(P^{-1}AP)P^{-1}\mathbf{X}P^{-T} + P^{-1}\mathbf{X}P^{-T}(P^TBP^{-T}) = P^{-1}CP^{-T}.$$

Unfortunately, this transformation yields coefficient matrices that are similar to the original ones, thus making the whole procedure useless. This simple example shows that different acceleration strategies need to be developed for the Sylvester equation; research has thus focused on constructing information-rich approximation spaces by using spectral transformations, rather than preconditioning, the way it is done in eigenvalue computations.

⁹One could also consider the existence of two matrices, one for A and one for B^T , respectively.

4.4.1. Projection methods. When both n and m are large, the dense solution matrix \mathbf{X} of (4.1) cannot be stored, therefore the determination of a memory saving approximation becomes mandatory. Whenever $C = C_1 C_2^\top$ has low rank, the results discussed in section 4.1 suggest that a low rank approximate solution may be determined, so that projection strategies are very appealing. Indeed, these methods compute low rank approximations $\tilde{\mathbf{X}} = V_k \mathbf{Y} W_j^\top \approx \mathbf{X}$, with V_k and W_j having far fewer columns than n and m , respectively, and are generalizations of the procedure seen in section 4.3.

Let \mathcal{V} and \mathcal{W} be two subspaces of \mathbb{C}^n , in principle not necessarily of the same dimension, and let the k (j) columns of V_k (of W_j) be orthonormal bases for \mathcal{V} (for \mathcal{W}), with $k \ll n$, $j \ll m$, such that \mathcal{V} is not orthogonal to $\text{range}(C_1)$ and \mathcal{W} is not orthogonal to $\text{range}(C_2)$. We look for an approximation $\tilde{\mathbf{X}} = V_k \mathbf{Y} W_j^\top \approx \mathbf{X}$, and we let $R := C_1 C_2^\top - A \tilde{\mathbf{X}} - \tilde{\mathbf{X}} B$ be the associated residual. Then we have $\tilde{\mathbf{x}} = \text{vec}(\tilde{\mathbf{X}}) = (W_j \otimes V_k) \text{vec}(\mathbf{Y})$, where $\tilde{\mathbf{x}}$ is an approximate solution of (4.3). Imposing a Galerkin (orthogonality) condition to the vector residual $c - A \tilde{\mathbf{x}}$ with respect to the space spanned by $W_j \otimes V_k$ corresponds to writing

$$(W_j \otimes V_k)^\top (c - A \tilde{\mathbf{x}}) = 0 \quad \Leftrightarrow \quad V_k^\top R W_j = 0. \quad (4.23)$$

Other conditions could be considered, such as the minimization of the residual in some norm, or the orthogonality of the residual with respect to some other space; see, e.g., [131], [127], [176].

If the columns of V_k and W_j span the spaces $K_k^\square(A, C_1)$ and $K_k^\square(B^\top, C_2)$, respectively, as in (4.16), then the obtained approximate solution $\tilde{\mathbf{X}} = V_k \mathbf{Y} W_j^\top$ may also be written as

$$\tilde{\mathbf{X}} = [C_1, A C_1, \dots, A^{k-1} C_1] G [C, B^\top C_2, \dots, (B^\top)^{k-1} C_2]^\top,$$

showing that projection methods yield a *polynomial* approximation to \mathbf{X} , which may be viewed as particular truncations of the finite sum closed form of the solution in (4.6) [223].

Substituting the residual matrix in the equation $V_k^\top R W_j = 0$ gives the following small size Sylvester equation:

$$V_k^\top A V_k \mathbf{Y} + \mathbf{Y} W_j^\top B W_j = V_k^\top C_1 (W_j^\top C_2)^\top. \quad (4.24)$$

If $V_k^\top A V_k$ and $W_j^\top B W_j$ have disjoint spectra, then this equation admits a unique solution for any right-hand side. By assuming that the fields of values of A and B are disjoint, one can ensure that $V_k^\top A V_k$ and $W_j^\top B W_j$ have disjoint spectra. Though restrictive, such an assumption is welcome also for stability purposes, so as to monitor that the solution \mathbf{X} have moderate norm. A typical implementation which simultaneously proceeds with both spaces is depicted in Algorithm 5.

Algorithm 5. Given A, B, C_1, C_2

1. Orthogonalize columns of C_1 to get $v_1 = V_1$
2. Orthogonalize columns of C_2 to get $w_1 = W_1$
3. For $k = 1, 2, \dots$
4. Compute \mathbf{Y}_k , solution to $(V_k^\top A V_k) \mathbf{Y} + \mathbf{Y} (W_k^\top B W_k) - V_k^\top C_1 C_2^\top W_k = 0$
5. If converged return V_k, \mathbf{Y}_k, W_k s.t. $\mathbf{X}_k = V_k \mathbf{Y}_k W_k^\top$ and stop
6. Compute next bases block:
 - Compute \hat{v} and \hat{w} for the chosen approximate space

Make \hat{v} orthogonal wrto $\{v_1, \dots, v_k\}$ and \hat{w} orthogonal wrto $\{w_1, \dots, w_k\}$
 Orthogonalize columns of \hat{v} to get v_{k+1} and columns of \hat{w} to get w_{k+1}
 Update: $V_{k+1} = [V_k, v_{k+1}]$, $W_{k+1} = [W_k, w_{k+1}]$

The process outlined in Algorithm 5 is very similar to that of Algorithm 4, the only difference being that here the space for B^\top also needs to be generated. For C_1, C_2 having p columns, and n and m rows, respectively (with, say, $m > n$), the computational cost at each iteration k can be summarized as follows:

- i) Solution of the projected problem: $\mathcal{O}((kp)^3)$ flops (see section 4.2);
- ii) Orthogonalization of the new basis vectors wrto the older vectors: $\mathcal{O}(mkp^2)$;
- iii) Orthogonalization of the new block, $\mathcal{O}(mp^2)$.

We also recall that in case the generated basis experiences loss of rank, standard deflation procedures can be applied to remove redundant columns, ensuring the reduction in the number of columns of the current basis block in subsequent iterations. Loss of rank may occur independently of the presence of an invariant subspace of the coefficient matrix, but simply due to the redundancy of some of the generated information; see [111] for a discussion in the context of multiple right-hand sides linear systems.

The computational cost of generating the next basis vectors \hat{v} and \hat{w} and the quality of the approximation both depend on the choice of \mathcal{V} and \mathcal{W} . This choice is usually based on similar arguments for each of the two spaces. We thus discuss the choice of \mathcal{V} , while the choice of \mathcal{W} can be made analogously. In his seminal article [211], Saad proposed Krylov subspaces for determining a low rank approximate solution to the Lyapunov equation by projection (the extension to the Sylvester equation is straightforward); the motivation was that Krylov subspaces tend to approximate well the action of the matrix exponential to a vector, so that the solution in the integral form (4.5) can take advantage of this property (see also section 5 for an explicit derivation). A major problem with this approach is that both bases V_k and W_j need to be stored to compute the final approximate solution. Since both matrices are full, this provides a severe limitation on the maximum affordable size of the two Krylov subspaces when A and B are large. In the quest of small but more effective spaces, several alternatives have been investigated. The impressive performance results of these enriched spaces have led to a resurgence of projection-type methods for linear matrix equations. In addition to the standard Krylov subspace, we list here a couple of recently explored selections for \mathcal{V} with A and C_1 ; similar choices can be made for \mathcal{W} using B^\top and C_2 .

- (a) Standard (block) Krylov subspace:

$$\mathcal{V} = \text{range}([C_1, AC_1, A^2C_1, \dots]);$$

- (b) Rational (block) Krylov subspace:

$$\mathcal{V} = \text{range}([(A + \sigma_1 I)^{-1}C_1, (A + \sigma_2 I)^{-1}(A + \sigma_1 I)^{-1}C_1, \dots]),$$

for a specifically chosen sequence $\{\sigma_j\}$, $j = 1, 2, \dots$ that ensure nonsingularity of the shifted matrix.

- (c) Global Krylov subspace:

$$\mathcal{V} = \left\{ \sum_{i \geq 0} A^i C_1 \gamma_i, \gamma_i \in \mathbb{R} \right\} = \text{span}\{C_1, AC_1, A^2C_1, \dots\},$$

where the linear combination is performed block-wise.

In all instances the least number of powers is computed so as to reach the dimension k . The subspaces listed above are somewhat related. For instance, the standard Krylov subspace can be formally obtained from the rational Krylov subspace for $\sigma_j = \infty$ for all j . Moreover, the Rational block Krylov subspace also includes the special choice of fixed poles at zero and infinity, which corresponds to the *extended* Krylov subspace in (4.21), namely $K_j^\square(A, C_1) + K_k^\square(A^{-1}, A^{-1}C_1)$, where j and k can in principle be different [72]. In addition, one can impose that C_1 belongs to the rational Krylov subspace with the choice $\sigma_1 = \infty$. The global Krylov subspace in (c) is a subspace of the block Krylov subspace; it was first proposed to solve linear systems with multiple right-hand sides [139], and then adapted to the Sylvester equation in [137]. Global spaces may be viewed as simplified versions of block Krylov spaces, where the polynomial coefficients are chosen to be multiples of the identity matrix, therefore lowering the number of degrees of freedom.

The criterion for stopping the iterative procedure in Algorithm 5, and thus the approximation spaces expansion, is usually based on the Frobenius or 2-norm of the residual matrix $R = A\bar{X} + \bar{X}B - C_1C_2^\top$. In general, R is dense and should not be computed explicitly if it has large dimensions. Its norm can be computed more cheaply if the generated spaces satisfy certain relations. Hence, assume that $\hat{v}_k, \hat{w}_j, \underline{H}_k$, and \underline{K}_j exist such that $AV_k = [V_k, \hat{v}_k]\underline{H}_k$, and $B^\top W_j = [W_j, \hat{w}_j]\underline{K}_j$, where $[V_k, \hat{v}_k]$ and $[W_j, \hat{w}_j]$ have orthonormal columns. If C_1 and C_2 satisfy $C_1 = [V_k, \hat{v}_k]C_1^{(k)}$, $C_2 = [W_j, \hat{w}_j]C_2^{(j)}$ for some $C_1^{(k)}, C_2^{(j)}$, then

$$\begin{aligned} \|R\|_F &= \|AV_k \mathbf{Y} W_j^\top + V_k \mathbf{Y} W_j^\top B - \hat{V}_k C_1^{(k)} (\hat{W}_j C_2^{(j)})^\top\|_F \\ &= \|[V_k, \hat{v}_k](\underline{H}_k Y[I, 0] + [I; 0] Y \underline{K}_j^\top - C_1^{(k)} (C_2^{(j)})^\top) [W_j, \hat{w}_j]^\top\|_F \\ &= \|\underline{H}_k Y[I, 0] + [I; 0] Y \underline{K}_j^\top - C_1^{(k)} (C_2^{(j)})^\top\|_F. \end{aligned} \quad (4.25)$$

The last expression involves a small matrix if k and j are small, and thus its norm can be cheaply evaluated. The spaces (a) to (c) above do satisfy the required conditions, and thus the residual norm can be monitored as the iteration proceeds.

All spaces listed above are nested, so that an approximate solution can be derived while each of these spaces is expanded.

The implementation can allow for different space dimensions for A and B , respectively, especially if the two coefficient matrices have rather different spectral properties. The idea of generating different approximation spaces - of the same dimension - for A and B by means of standard Krylov subspaces was first developed in [131], where however the right-hand side C of the original problem was approximated by a rank-one matrix $c_1 c_2^\top$, to be able to build the standard Krylov subspaces $K_j(A, c_1)$ and $K_j(B^\top, c_2)$ as approximation spaces. The approach was then generalized to block Krylov subspaces in [223], so as to exploit the low (but possibly larger than one) rank matrices C_1, C_2 . Distinct Krylov subspaces for the right and left subspaces should be considered also when $B = A^\top$, as long as $C_1 C_2^\top$ is nonsymmetric. Nonetheless, in this case the generation of the two spaces can share some computationally intensive work, such as shifted system solves with the same coefficient matrix. The possibility of using nonsymmetric Lanczos processes which simultaneously generate $K_j(A, C_1)$ and $K_j(A, C_2)$ was explored in [134].

In Figure 4.1 we report a typical convergence history for the norm of the residual matrix, when the standard Krylov and extended Krylov subspaces are used for both A and B . Here data refer to A as the finite difference discretization of the Laplace operator in the unit square with homogeneous boundary conditions, and B as the

same type of discretization for the operator $\mathcal{L}u = (\exp(-4xy)u_x)_x + (\exp(4xy)u_y)_y$, leading to matrices of the same size, $40,000 \times 40,000$; $C = c_1 c_2^\top$ with c_1, c_2 vectors with normally distributed random entries. We notice that with a subspace of dimension less than 140 for each matrix, the extended Krylov subspace is able to reduce the norm of the residual matrix by ten orders of magnitude, whereas the standard Krylov subspace of the same dimension shows a very modest residual norm reduction. In terms of computational costs of the extended procedure, the matrices A and B are pretty sparse and systems with them can be efficiently solved by a sparse direct solver.

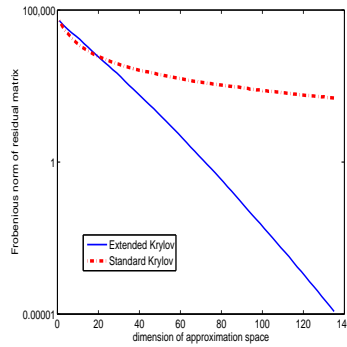


FIG. 4.1. *Typical convergence history of projection methods for the Sylvester equation with standard and Extended Krylov subspaces.*

Rational Krylov subspaces have a rich history. First introduced by Ruhe in the context of eigenvalue approximation [210], their relevance has significantly spread in applied approximation theory and model order reduction frameworks, due to their functional approximation properties; see, e.g., [4],[106],[112] and references therein.

The effectiveness of general rational spaces strongly relies on the efficiency of solving systems with A or its shifted variants. The reliability of recent direct sparse and iterative linear system solvers has made it possible to use these richer approximation spaces for more complex problems like the ones we are addressing. The choice of the shift is crucial to achieve fast convergence; this issue is postponed to the corresponding discussion for the Lyapunov equation in section 5.2.1.

In the quest for memory savings, the possibility of restarting the process could be considered: a maximum subspace dimension is allowed and the final approximate solution is obtained as $\tilde{\mathbf{X}} = \tilde{\mathbf{X}}^{(0)} + \tilde{\mathbf{X}}^{(1)} + \tilde{\mathbf{X}}^{(2)} + \dots$, where the superscripts indicate a new restart. Strategies on how to generate the new approximations were proposed in [131]. We mention that new restarting procedures were recently proposed in [3], but their overall computational costs for large scale matrices have not clearly been assessed. An alternative that could be considered in the symmetric case is to resort to a two-pass strategy, inspired by a similar procedure in the eigenvalue context. Indeed, for A and B symmetric, not necessarily equal, an orthogonal basis of each standard Krylov subspace together with the projected matrix could be generated without storing the whole basis, but only the last three (block) vectors, because the orthogonalization process reduces to the short-term Lanczos recurrence [212]. Therefore, in a first-pass only the projected solution \mathbf{Y} could be determined while limiting the storage for V_k and W_j ; at convergence the factors of the approximate solution $\tilde{\mathbf{X}} = V_k \mathbf{Y} W_j^\top$ could be recovered by generating the two bases once again; an

implementation of such an approach can be found in [158] for $B = A^\top$ and $C_1 = C_2$. The same idea could be used for other situations where a short-term recurrence is viable; the effectiveness of the overall method strongly depends on the affordability of computing the two bases twice.

The convergence analysis of projection methods has long been overlooked. Following recent significant advances in the convergence study of projection methods for the Lyapunov equation (see section 5.2.1) Beckermann in [21] provided a thorough study: residual norm bounds are given for Galerkin projection methods when Rational Krylov subspaces, of possibly different dimensions, are used for A and B^\top . The proposed estimates rely on new residual matrix relations, and highlight the role of the field of values of the two coefficient matrices; we refer the reader to Proposition 4.4 below for further details on the results in [21]. Advances in the theoretical aspects of projection methods have been made in close connection with the recent great steps ahead taken in the understanding of polynomial and rational approximation methods for matrix functions such as the matrix exponential. The interplay of numerical linear algebra, approximation theory and functional analysis has made this possible; see, e.g., [124],[112] and their references.

4.4.2. ADI iteration. The Alternating-Direction-Implicit (ADI) iteration was first introduced in [192] in 1955, and proposed to solve large Sylvester equations by Ellner and Wachspress in [78]¹⁰. Since then, and with various computationally effective refinements, the approach has been one of the leading methods for solving large-scale Sylvester (and Lyapunov) equations. In its original form discussed in [78] and summarized next, the ADI iteration is derived for a full matrix \mathbf{X} (see also Smith [231] for the derivation below). A low memory factorized version is used in practice for large matrices, and it will be presented in the sequel. In the following we assume that both real matrices A and B have eigenvalues with positive real parts. We can equivalently rewrite (4.1) as

$$(qI + A)\mathbf{X}(qI + B) - (qI - A)\mathbf{X}(qI - B) = 2qC, \quad q \neq 0.$$

For $q > 0$, $qI + A$ and $qI + B$ are nonsingular and we can multiply by their inverses so as to obtain the following equation:

$$\mathbf{X} - (qI + A)^{-1}(qI - A)\mathbf{X}(qI - B)(qI + B)^{-1} = 2q(qI + A)^{-1}C(qI + B)^{-1}.$$

Let $\mathcal{A} = (qI + A)^{-1}(qI - A)$, $\mathcal{B} = (qI - B)(qI + B)^{-1}$ and $\mathcal{C} = 2q(qI + A)^{-1}C(qI + B)^{-1}$. With this notation, the matrix equation above has the form $\mathbf{X} - \mathcal{A}\mathbf{X}\mathcal{B} = \mathcal{C}$, and it is called the Stein equation ; see section 6. The matrix $\mathbf{X} = \sum_{k=1}^{\infty} \mathcal{A}^{k-1}\mathcal{C}\mathcal{B}^{k-1}$ is a formal solution to the Stein equation, and since both \mathcal{A} and \mathcal{B} have spectral radius less than one¹¹, the series is convergent. This consideration drives the implementation of the following sequence of approximations

$$\mathbf{X}_0 = \mathcal{C}, \quad \mathbf{X}_{k+1} = \mathcal{C} + \mathcal{A}\mathbf{X}_k\mathcal{B}. \quad (4.26)$$

The approach can be generalized to two parameters $p, q > 0$ for A and B , respectively, giving the transformed equation

$$\mathbf{X} - \mathcal{A}(p, q)\mathbf{X}\mathcal{B}(p, q) = \mathcal{C}(p, q),$$

¹⁰The authors of [78] referred to these Sylvester equations as Lyapunov equations.

¹¹For a given matrix A with eigenvalues λ in \mathbb{C}^+ and $q > 0$, the eigenvalues of $(qI - A)(qI + A)^{-1}$ are given by $(q - \lambda)/(q + \lambda)$, with absolute values all less than one.

with $\mathcal{A}(p, q) = (pI + A)^{-1}(A - qI)$, $\mathcal{B}(p, q) = (B - pI)(qI + B)^{-1}$ and $\mathcal{C}(p, q) = (p + q)(pI + A)^{-1}C(qI + B)^{-1}$. A recursion similar to the one for a single parameter can be derived, and it is convergent if the spectral radii of $\mathcal{A}(p, q), \mathcal{B}(p, q)$ are both less than one. Therefore, the parameters p, q are selected so as to minimize these spectral radii, and if A, B are both symmetric with spectral intervals (a, b) and (c, d) , respectively, this corresponds to solving the ADI minimax problem

$$\min_{p, q > 0} \max_{s \in (a, b), t \in (c, d)} \left| \frac{(q - s)(p - t)}{(p + s)(q + t)} \right|.$$

The generalization of this concept allows one to choose different p, q at each iteration, providing a sequence of parameters p_1, p_2, \dots and q_1, q_2, \dots . The associated ADI minimax problem after J iterations thus becomes

$$\min_{p_j, q_j > 0} \max_{s \in (a, b), t \in (c, d)} \prod_{j=1}^J \left| \frac{(q_j - s)(p_j - t)}{(p_j + s)(q_j + t)} \right|,$$

which, if solved exactly, provides optimal parameters for the convergence rate of the ADI iteration. In practice, a fixed number J of parameters is selected a priori and then cyclically repeated until convergence. The choice of J is driven by the quality of the computed parameters: few parameters may be better than many badly distributed parameters. We will return to this issue in section 5.2.

Following a successful idea developed for the Lyapunov equation, the authors of [25] propose a factorized version of the ADI iteration, which allows one to write the approximate solution as the product of three memory saving factors, as long as $C = C_1 C_2^\top$ is low rank. We will expand on this and other implementation aspects in the case of the Lyapunov equation, since the changes occurring when generalizing ADI to the Sylvester equation are mainly technical and are due to the presence of a left space and a right space; we point here to the recent works of Peter Benner and his collaborators for a comprehensive implementation investigation of ADI for the Sylvester equation.

We conclude with a theoretical comparison recently made between ADI and the Galerkin method (see section 4.4.1) in the Rational Krylov subspaces

$$K_m(A, C_1) = \text{range}\{C_1, (A + \sigma_2 I)^{-1} C_1, \dots, (A + \sigma_m I)^{-1} C_1\},$$

$$K_m(B^\top, C_2) = \text{range}\{C_2, (B^\top + \eta_2 I)^{-1} C_2, \dots, (B^\top + \eta_m I)^{-1} C_2\},$$

where in both cases the first pole is taken to be at infinity, so that the columns of C_1 and C_2 belong to the corresponding spaces. In [86], Flagg and Gugercin showed that ADI and the Galerkin approach are equivalent whenever the poles of both methods coincide with the eigenvalues of the projections of A and B (Ritz values) in the two spaces, respectively; the same result was earlier proved for the Lyapunov equation with different techniques (see Theorem 5.6). Moreover, for general poles the following result is proved by Beckermann for the error [21, Corollary 2.2].

PROPOSITION 4.4. *Let \mathbf{X} be the exact solution to the Sylvester equation. Let $\mathcal{S}_{A,B}X = AX + XB$, and let $\mathbf{X}_{m,m}^G, \mathbf{X}_m^{ADI}$ be the approximate solutions obtained after m iterations of the Galerkin method in $K_m(A, C_1), K_m(B^\top, C_2)$, and after m ADI steps, respectively, with the two methods using the same poles. If the fields of values $W(A)$ and $W(-B)$ have empty intersection, then*

$$\|\mathcal{S}_{A,B}(X - X_{m,m}^G)\|_F \leq \gamma_0 \|\mathcal{S}_{A,B}(X - X_m^{ADI})\|_F$$

with a constant $\gamma_0 \leq 3 + 2c_0$, with $c_0 = 2\text{diam}(W(A), W(-B))/\text{dist}(W(A), W(-B))$ independent of the poles used to generate the space.

The constant γ_0 is not optimal. As stated in [21], Proposition 4.4 shows that, even for optimal poles, ADI cannot give much better results than rational Galerkin; moreover, for poor poles ADI is known to give much larger residuals. Further results will be discussed for the case of the Lyapunov equation.

4.4.3. Data sparse and other methods. A variety of approaches relying on the data sparsity structure has been analyzed. These methods may be particularly appropriate in the large scale case when the right-hand side matrix C is sparse and full rank.

The Kronecker formulation allows one to consider a wide range of linear system solvers for (4.3); an early ad-hoc implementation of the classical SOR was proposed in [235], although the exploding dimensions of the problem significantly penalize the method, when compared with the approaches analyzed so far. We also recall from section 4.4.1 that global Krylov subspace methods represent an implicit way to deal with the Kronecker formulation. Other iterative solvers based on the Kronecker formulation (4.3) have been explored specifically for the Lyapunov equation, and they will be reviewed in section 5.2.3. These appear to be the main directions taken whenever C is not numerically low rank.

For data-sparse matrices $A \in \mathbb{R}^{n \times n}$ and $B \in \mathbb{R}^{m \times m}$, namely such that matrix-vector multiplications for A and B can be performed with complexity $\mathcal{O}(n)$ and $\mathcal{O}(m)$, respectively, an extension of the multigrid algorithm was more recently proposed in [101], where A and B stem from the discretization of a class of partial differential equations, and their spectra are assumed to be separated by a line. A particular computational caveat of this extension is the smoother (e.g., Jacobi), which in this case requires approximately solving a *diagonal* Sylvester equation at each iteration. This step is carefully discussed in [101], and a procedure for determining a cheap and low rank approximate solution is devised. Other crucial points include handling the connection between the (independently generated) sequences of matrices for A and B , which is accounted for during the smoothing procedure, and the imposition of regularity constraints on the continuous operators associated with A and B . A major issue arising when using these hierarchical methods is whether the approximate solution $\tilde{\mathbf{X}}$ is low rank, so that it can be stored cheaply by means of a (hierarchical) sparse format, the \mathcal{H} -matrix format. Such a format is a data-sparse representation for a special class of matrices, which appear to occur after the use of several discretization methods, when PDEs or integral equations are treated numerically [103]. The \mathcal{H} -matrix format consists of partitioning a given matrix recursively, into submatrices admitting low-rank approximations. The definition of this format requires the introduction of further arithmetic operations/approximations, so as to be able to determine, e.g., an \mathcal{H} -matrix after the approximate inversion of an \mathcal{H} -matrix, so as to make the class closed with respect to some important matrix operations; see section 5.2.3 for further details.

A different though related approach consists in adapting small scale iterations to the large setting, again under the condition that C is low rank. This can be performed, for instance, within the sign function iteration, by using rank truncation of the iterates, and sparse format for the approximate solution. More details on the sign function iteration will be given in section 5.2.3. Here we mention that such an approach is investigated in [18] (see also [20]), where the sparse format chosen for the data and for the approximate solution is the hierarchical \mathcal{H} -matrix format also used in

[99], [101]. With this approach, sparse approximate solutions to a Sylvester equation of size up to $n = 262144$ associated with a control problem for the two-dimensional heat equation are reported in [18]. The accuracy and effectiveness of the method depend on some thresholds used for maintaining sparsity and low rank during the iteration, and are thus problem dependent.

5. Continuous-time Lyapunov equation. For $B = A^\top$ from the Sylvester equation we obtain the Lyapunov equation

$$A\mathbf{X} + \mathbf{X}A^\top = C, \quad (5.1)$$

with C symmetric, and its generalized counterpart, $A\mathbf{X}E^\top + E\mathbf{X}A^\top = C$, with E nonsingular. Clearly, this latter equation can be transformed into the form (5.1) by left and right multiplication by E^{-1} and $E^{-\top}$, respectively. If E is symmetric and positive definite, a Cholesky decomposition could be performed and its inverse factors applied to the equation on the left and right sides, so as to maintain the problem structure. These are called the *continuous-time* Lyapunov equations, to be distinguished from the discrete-time equations, which will be discussed in section 6. They arise in the analysis of continuous-time and discrete-time linear dynamical systems, respectively. A very detailed analysis of the Lyapunov equation, with computational developments until 1995 and many relevant connections in the control application area can be found in [88].

In the context of inertia theory, (5.1) with $C \succeq 0$ relates the location of the eigenvalues of both A and \mathbf{X} with respect to the imaginary axis. Since C is symmetric, the solution \mathbf{X} is also symmetric. According to the Sylvester equation theory, the solution to (5.1) exists and is unique if and only if $\lambda_i + \bar{\lambda}_j \neq 0$ for all eigenvalues λ_i, λ_j of A [130]. If all eigenvalues of A have negative real part, namely A is stable, then this condition is satisfied, so that a unique solution is ensured. We remark that the stability of A is an important property in the control setting, therefore it is not regarded as a restriction for solving the Lyapunov equation, although not strictly required. We shall see, however, that some of the large-scale methods require additional restrictions on A , namely its negative definiteness, to ensure the existence of an approximate solution. For A nonsymmetric, this extra condition may limit the applicability of the method, since in general a stable matrix A is not necessarily negative definite.

It can be verified that if A is stable and $C \succ 0$ ($C \succeq 0$) then $\mathbf{X} \succ 0$ ($\mathbf{X} \succeq 0$); in this case the problem is called the *stable* Lyapunov equation. If $C \succeq 0$ and (A, C^\top) is observable, then $\mathbf{X} \succ 0$. A detailed account of various relations between the inertia of A and of \mathbf{X} can be found, e.g., in [165, section 13.1], [219],[220]. A specialized sensitivity bound can be obtained for the stable Lyapunov equation. Assume that $\mathbf{X} + \Delta\mathbf{X}$ solves

$$(A + \Delta A)(\mathbf{X} + \Delta\mathbf{X}) + (\mathbf{X} + \Delta\mathbf{X})(A + \Delta A)^\top + (C + \Delta C) = 0,$$

then

$$\frac{\|\Delta\mathbf{X}\|}{\|\mathbf{X} + \Delta\mathbf{X}\|} \leq 2\|A + \Delta A\| \|H\| \left[\frac{\|\Delta A\|}{\|A + \Delta A\|} + \frac{\|\Delta C\|}{\|C + \Delta C\|} \right],$$

where H satisfies $AH + HA^\top + I = 0$, and all denominators are assumed to be nonzero [120]. Estimates for the backward error associated with the Lyapunov equation do not differ from those in (4.10) for the Sylvester equation; therefore, except for the substitution $B = A^\top$, the extra structure of the problem does not modify the sensitivity properties of the solution [122].

The sensitivity of the solution to (5.1) can also be analyzed by looking at the spectral properties of the solution matrix; this topic has attracted a lot of interest, especially in light of its consequences in the stability analysis of dynamical systems. Various authors have explored the spectral decomposition of the Lyapunov solution, to make statements on the decay of its eigenvalues; see, e.g., [197],[232],[5],[154]. In [5], an error estimate for a low rank approximation to the solution of (5.1) was proved. For the sake of simplicity we report here only the case when C is rank-one. The result relies on the fact that the solution matrix admits the following decomposition:

$$\mathbf{X} = ZDZ^\top, \quad D = \text{diag}(\delta_1, \dots, \delta_n), \quad \delta_k = \frac{-1}{2\Re(\lambda_k)} \prod_{j=1}^{k-1} \left| \frac{\lambda_k - \lambda_j}{\lambda_k^* + \lambda_j} \right|^2,$$

where λ_j are the eigenvalues of the diagonalizable matrix A .

THEOREM 5.1. *Assume A is diagonalizable with eigenvector matrix Q , having all unit norm columns, and let $C = cc^\top$. Let $\mathbf{X} = \sum_{j=1}^n \delta_j z_j z_j^\top$ solve (5.1), with the nonnegative values δ_j sorted decreasingly, and for $k \in \{1, \dots, n\}$ define $\mathbf{X}_k = \sum_{j=1}^k \delta_j z_j z_j^\top$. Then*

$$\|\mathbf{X} - \mathbf{X}_k\| \leq (n - k)^2 \delta_{k+1} (\kappa(Q) \|c\|_2)^2,$$

where $\|\cdot\|$ is the matrix norm induced by the vector 2-norm.

The bound may not be sharp for highly non-normal A , for which $\kappa(Q)$ may be large. A more specialized bound was earlier given by Penzl for A symmetric, which only depends on the condition number of A [197].

THEOREM 5.2. ([197]) *Let A be symmetric and negative definite, with condition number $\kappa(A)$, and $C = C_1 C_1^\top$ with C_1 of rank p . Let $\lambda_i(\mathbf{X})$ with $i = 1, \dots, n$ be the nonincreasingly ordered eigenvalues of \mathbf{X} . Then*

$$\frac{\lambda_{pk+1}(\mathbf{X})}{\lambda_1(\mathbf{X})} \leq \left(\prod_{j=0}^{k-1} \frac{\kappa^{(2j+1)/(2k)} - 1}{\kappa^{(2j+1)/(2k)} + 1} \right)^2, \quad \text{for } 1 \leq pk < n.$$

It was also shown in [197] that eigenvalues alone cannot predict the eigenvalue distribution of \mathbf{X} in the nonhermitian case; this is reminiscent of similar limitations of non-normal matrices in the convergence analysis of iterative methods for linear systems [174, sec. 5.7.3].

Bounds on the eigenvalue decay that attempt to cope with non-normality were obtained in [213, sec.3.1.2], where the concept of pseudospectrum is used; in there, some interesting counter-intuitive convergence behaviors are also described. Overall, much remains to be understood in the decay of the solution spectrum in the non-normal case.

In addition to the application relevance, establishing conditions under which the solution matrix has exponentially decaying eigenvalues provides theoretical ground for the good performance of low rank projection methods in the large scale case.

5.1. Lyapunov equation. Small scale computation. As for the Sylvester equation, the closed form solutions described in section 4 could be used in theory. A detailed account of early methods can be found in [88], together with some ad-hoc algorithms appropriate when special forms of A - e.g., Schwarz, Companion or Jordan forms - are available; see also [39], [119] for improved approaches for the companion form.

The standard method for efficiently solving (5.1) when A has small dimensions does not essentially differ from those for the Sylvester equation discussed in previous sections. In fact, since $B = A^\top$, the computational cost of the reduction to Schur form is halved in the Bartels-Stewart method [221].

A specifically designed algorithm was proposed by Hammarling, to exploit the case when C is positive semidefinite. It was shown in [113] that if $C = C_1 C_1^\top \succeq 0$, it is possible to determine the Cholesky factor L of the solution $\mathbf{X} = LL^\top$, without first determining \mathbf{X} . The computation of the Cholesky factor has some advantages when \mathbf{X} is nonsingular but severely ill-conditioned, so that dealing with L significantly improves the accuracy and robustness of computations with \mathbf{X} ; in [273] a comparison between Hammarling's and Bartels-Stewart methods can be found. A block variant of Hammarling's method for the discrete-time Lyapunov equation is suggested in [156], which can dramatically improve the performance of the original scalar (unpartitioned) algorithm on specific machine architectures, while preserving the stability of the original method.

We also mention the possibility of pre-processing, proposed both in the continuous and discrete-time equations, so as to transform the original symmetric problem onto a skew-symmetric one, so that the solution will also be skew-symmetric ($\mathbf{X} = -\mathbf{X}^\top$), allowing for some memory savings; see [88, sec. 2.1.2] and references therein.

A completely different approach exploits the fact that the solution \mathbf{X} may be computed by means of matrix functions, in particular, by using the sign function. Although less general than Schur-form-based algorithms, they allow one to handle larger problems, especially if the right-hand side is low-rank or structured, and can be more easily adapted to a high performance computational environment. The idea is to use well-established matrix iterations to obtain the matrix sign function in a cheap manner, by fully exploiting the possible sparse format of the matrix. The whole procedure is actually more general, and it also applies to the symmetric algebraic Riccati equation (see Algorithm 1). Here we will follow the derivation proposed in [19], see also [27], although the main iteration was introduced by Larin and Aliev in [167] for the generalized Lyapunov equation. Let $\mathcal{A} = \mathcal{X} \text{blkdiag}(J_+, J_-) \mathcal{X}^{-1}$ be the Jordan decomposition of a given matrix \mathcal{A} , where J_+, J_- represent the Jordan matrices associated with the eigenvalues in the open planes \mathbb{C}^+ and \mathbb{C}^- , respectively. Then $\text{sign}(\mathcal{A}) = \mathcal{X} \text{blkdiag}(I, -I) \mathcal{X}^{-1}$, where the dimensions of $I, -I$ match those of J_+ and J_- , respectively. For A stable, the solution to the Lyapunov equation satisfies (see, e.g., [208])

$$\begin{bmatrix} 0 & \mathbf{X} \\ 0 & I \end{bmatrix} = \frac{1}{2} \left(I + \text{sign} \left(\begin{bmatrix} A^\top & C \\ 0 & -A \end{bmatrix} \right) \right) =: \frac{1}{2} (I + \text{sign}(Z_0)). \quad (5.2)$$

With this property, the following matrix iteration corresponds to applying the Newton method to the nonlinear equation $(\text{sign} Z_0)^2 = I$:

$$Z_{k+1} = \frac{1}{2} (Z_k + Z_k^{-1}), \quad k = 0, 1, \dots, \quad (5.3)$$

yielding

$$\text{sign} Z_0 = \lim_{k \rightarrow \infty} Z_k = \begin{bmatrix} -I & 2\mathbf{X} \\ 0 & I \end{bmatrix}.$$

Although the iteration is globally and (asymptotically) quadratically convergent, the basic iteration above may have slow initial convergence, therefore it is often accelerated

by using a parameterized procedure, that is $Z_{k+1} = \frac{1}{2}(c_k Z_k + (c_k Z_k)^{-1})$, $k = 0, 1, \dots$, for an appropriate selection of the parameter $c_k > 0$. A popular choice is $c_k = |\det(Z_k)|^{-\frac{1}{n}}$ [50]; see [7] for a review of other choices.

5.2. Lyapunov equation. Large scale computation. Recalling the discussion for the Sylvester equation in section 4.4, the solution of the Lyapunov equation for A of large dimensions focuses on the determination of memory saving and computationally appealing approximations. For the stable problem, this is achieved in most cases by looking for a *low* rank approximation $\tilde{\mathbf{X}} = \mathbf{Z}\mathbf{Z}^*$, so that only the tall matrix \mathbf{Z} is actually computed and stored. This is possible if, for instance, the right-hand side has low rank, since in that case we also have $\mathbf{X} \succeq 0$. Nonetheless, strategies to approximate the general right-hand side by low-rank matrices have also been explored in the literature; see, e.g., [131].

To help fully grasp the relevance of the topic, we notice that a number of recent PhD theses have been devoted to the theory and computational aspects of the large-scale Lyapunov matrix equation, whose results have significantly advanced knowledge on the problem; among them, [195],[187],[183],[128],[273],[172],[213]. The list would expand if one were to also include closely related theses on model order reduction of linear dynamical systems.

We conclude this section by noticing that a systematic numerical comparison of all iterative methods described in the following subsections on a variety of very large problems (of size $n \gg 10^4$) is still lacking, although in our presentation some guidelines are given about the settings in which each of the discussed method should be preferred.

5.2.1. Projection methods. As in the case of the Sylvester equation, the derivation of a projection method can be determined by imposing, e.g., the Galerkin condition to the residual, with respect to some approximation space. In particular, from equation (4.24) with $k = j$, $V_k = W_j$ and $C_2 = C_1$, we obtain the following projected small size Lyapunov equation:

$$V_k^* A V_k \mathbf{Y}_k + \mathbf{Y}_k V_k^* A^\top V_k = V_k^* C_1 (V_k^* C_1)^*, \quad (5.4)$$

whose solution matrix \mathbf{Y}_k gives $\mathbf{X}_k = V_k \mathbf{Y}_k V_k^* \approx \mathbf{X}$. Since \mathbf{Y}_k is positive semi-definite and numerically singular, it is possible to perform a truncated decomposition of \mathbf{Y}_k as $\mathbf{Y}_k = LL^*$, so that only the slim factor $Z_k = V_k L$ of the solution $\mathbf{X}_k = Z_k Z_k^*$ needs to be stored. To ensure that (5.4) admits a unique solution, the matrix $V_k^* A V_k$ is assumed to be stable. Such a sufficient condition is met by requiring that A be negative definite, which is the usual hypothesis when using projection methods. This condition represents a limitation of projection methods, since the original problem admits a unique solution even in case of a stable¹² but not necessarily negative definite A . On the other hand, these are sufficient conditions: projection methods can work in practice without this assumption, although they may break down or show some erratic convergence behavior; see [176] for an analysis.

An apparently different (functional) approach, based on the approximation to the matrix exponential and on (4.4), leads exactly to the same approximation procedure as Galerkin. Indeed, the action of the matrix exponential to C_1 , $\exp(tA)C_1$, can be approximated in the space \mathcal{V} as $V_k \exp(tH_k)(V_k^* C_1)$ where $H_k = V_k^* A V_k$, so that the analytic expression in (4.4) for the solution can be approximated explicitly; this is the

¹²As already said, even stability of A is not strictly necessary for the solvability of the Lyapunov equation, but only that $I \otimes A + A \otimes I$ be nonsingular.

way the Galerkin approximate solution was originally obtained in [211] for a rank-one matrix C_1 .

PROPOSITION 5.3 ([211]). *Let \mathcal{V} be a subspace of \mathbb{R}^n , and let V be such that $\mathcal{V} = \text{range}(V)$. Let also H be the projection and restriction of A onto \mathcal{V} , and $y_m = \exp(tH)(V^*C_1)$. Then the matrix $V\mathbf{Y}V^*$ with*

$$\mathbf{Y} = \int_0^\infty y_m(t)y_m(t)^* dt$$

is the Galerkin approximate solution to the Lyapunov equation in \mathcal{V} .

The procedure above is very general, and the success of the approach, in terms of computational cost, depends on the choice of the approximation space \mathcal{V} . All choices discussed in section 4.3 have been explored. For instance, the block Krylov subspace $K_k^\square(A, C_1)$ was exploited in [135], and it was referred to as the Arnoldi method, after the procedure used to build the block Krylov subspace. The following computationally convenient relation for the residual $R_k = A\mathbf{X}_k + \mathbf{X}_kA^* + C_1C_1^*$ can be deduced from (4.25) ([135, Th. 2.1]):

$$\|R_k\|_F = \sqrt{2}\|(v_{k+1}^*Av_k)E_k^*\mathbf{Y}_k\|_F, \quad E_k^* = [0_m, \dots, 0_m, I_m],$$

where v_{k+1} contains the next block of basis vectors.

Finally, the solution \mathbf{X}_k is the exact solution to the nearby problem ([135])

$$(A - \Delta)\mathbf{X} + \mathbf{X}(A - \Delta)^* - C_1C_1^* = 0,$$

with $\Delta = V_{k+1}(V_{k+1}^*AV_k)V_k^*$, $\|\Delta\|_F = \|V_{k+1}^*AV_k\|_F$.

The asymptotic convergence of the Arnoldi method was recently analyzed in [226]. Here we report a sample of such analysis, which applies to A symmetric and positive definite, and C_1 of rank one and unit norm; the derived bound was shown in [226] to provide an accurate worst-case convergence rate of the method.

THEOREM 5.4. *Let A be symmetric and positive definite, and let λ_{\min} be the smallest eigenvalue of A . Let $\hat{\lambda}_{\min}$, $\hat{\lambda}_{\max}$ be the extreme eigenvalues of $A + \lambda_{\min}I$ and $\hat{\kappa} = \hat{\lambda}_{\max}/\hat{\lambda}_{\min}$. Let \mathbf{X}_m be the Galerkin approximate solution to \mathbf{X} in a Krylov subspace of dimension m . Then*

$$\|\mathbf{X} - \mathbf{X}_m\| \leq \frac{\sqrt{\hat{\kappa}} + 1}{\hat{\lambda}_{\min}\sqrt{\hat{\kappa}}} \left(\frac{\sqrt{\hat{\kappa}} - 1}{\sqrt{\hat{\kappa}} + 1} \right)^m, \quad (5.5)$$

where the matrix norm is the one induced by the vector 2-norm.

This result shows that the error norm is bounded by the same asymptotic quantity as for the Conjugate Gradient method applied to a standard linear system with coefficient matrix $A + \lambda_{\min}I$.

As already mentioned, the algorithmic steps to compute an approximate solution by projection remain unchanged when a different approximation space is used. In [224] an efficient method based on the extended Krylov subspace in (4.21), namely $K_k^\square(A, C_1) + K_k^\square(A^{-1}, A^{-1}C_1)$, was introduced for C_1 of low rank. In [224] the method was experimentally shown to be highly superior to the Arnoldi method for sparse and large coefficient matrices, allowing the computation of an equally accurate solution with a significantly smaller dimensional subspace, at lower computation costs for sparse A ; convergence plots are typically similar to that in Figure 4.1. According to the experiments in [224], the method also compares well with respect to ADI. A

recent asymptotic analysis in [152] theoretically confirmed these results, showing that the Extended Krylov subspace method is expected to have higher convergence rate than the Arnoldi method, and than ADI with a single pole.

More recently, rational Krylov subspaces have been introduced as approximation space in the solution of the Lyapunov equation, showing the great potential of projection methods, when solving shifted systems with A is affordable [73]. In [76] estimates for the error norm were derived, assuming that the field of values $W(A)$ of A and the set of parameters $\{\sigma_i\}_{i=1,2,\dots}$ are both contained in the same half plane. We first report on an error bound that emphasizes the connection with ADI.

THEOREM 5.5. *Let $C_1 \in \mathbb{R}^n$, and let \mathbf{X}_m be the approximation obtained by a Galerkin method in range($(A + \sigma_1 I)^{-1}C_1, \dots, (A + \sigma_m I)^{-1}C_1$). Let $\gamma_1 > 0$ be the Crouzeix constant (with $\gamma_1 \leq 11.08$) and $\gamma_2 = \frac{1}{2\text{dist}(0, W(A))}$. Then*

$$\|\mathbf{X} - \mathbf{X}_m\|_F \leq 2\gamma_1^2 \gamma_2 \max_{z \in W(A)} \frac{\prod_i |z - \bar{\sigma}_i|^2}{\prod_i |z + \sigma_i|^2} \|C_1\|^2, \quad (5.6)$$

where $\|\cdot\|_F$ is the Frobenius norm.

The maximization problem appearing in the upper bound is the same as that characterizing the convergence rate in the ADI method (see section 5.2.2). We also note that the ADI bound is reachable, in terms of convergence rate, therefore it may be viewed as a worst case scenario for ADI. Therefore, the result of Theorem 5.5 provides a realistic picture of the performance of the Rational Krylov subspace method compared with that of ADI, whose implementation for the Lyapunov equation is described in detail in section 5.2.2. The superiority of the rational Krylov subspace method for the same shifts can be easily appreciated in practical cases; we refer the reader to, for instance, Examples 4.3-4.4 and Fig. 4.5 in [76]. Beckermann in [21] expanded this type of result to the more general setting of the Sylvester problem (see Proposition 4.4). Nonetheless, if close to optimal poles can be selected, ADI may still provide a competitive alternative to rational Krylov subspace methods.

Asymptotic error bounds for the Galerkin method in the Rational Krylov subspace were derived in [76]. The reported numerical experiments on worst-case spectral distributions show that these bounds are indeed sharp for certain classes of data.

The general rational Krylov subspace requires the selection of a series of shifts (poles), which can either be computed a-priori or during the generation of the space. The a-priori pole computation may require possibly high computational costs, following the same procedure as that used for other parameter-dependent methods such as ADI (see below). On the other hand, it was shown in [74] that it is possible to employ a greedy algorithm to compute the next pole on the fly, while the iteration proceeds, with computational costs of a modest power of the order of the space dimension, which is usually significantly smaller than the problem dimension. This is done by exploiting approximate spectral information generated within the current approximation space. More precisely, for $C_1 = c_1 \in \mathbb{R}^n$, we first observe that an element of the rational Krylov subspace of dimension m can be written as $x = p_{m-1}(A)q_m(A)^{-1}c_1$ with p_{m-1} and q_m polynomials of degree $m-1$ and m , respectively, and where the roots of q_m are the parameters $\sigma_1, \dots, \sigma_m$. For the sake of the derivation, assume that the linear system $(A + sI)x = c_1$ for some parameter $s > 0$ is to be solved, then the residual of an approximate solution x_m obtained by imposing the Galerkin condition

with respect to the space can be written as

$$c_1 - (A + sI)x_m = \frac{r_m(A)c_1}{r_m(s)}, \quad r_m(z) = \prod_{j=1}^m \frac{z + \lambda_j}{z + \sigma_j}$$

where $\lambda_1, \dots, \lambda_j$ are the eigenvalues of the projection of A onto the space (Ritz values). The adaptive procedure amounts to determining the next parameter σ_{m+1} by capturing the parameter σ for which the current rational function is largest:

$$\sigma_{m+1} = \arg \left(\max_{\sigma \in \partial S_m} \frac{1}{|r_m(\sigma)|} \right),$$

where $S_m \subset \mathbb{C}^+$ approximates the spectral region of A , and ∂S_m is its border. The actual computational procedure requires an initial rough estimate of ∂S_m , which for real A can be taken to be the approximate extreme values of the interval $\partial S_m \cap \mathbb{R}$. These can be easily obtained by few iterations of an eigenvalue solver [74].

Numerical experiments reported in [74] show that the method is superior to the Extended Krylov subspace when, for instance, the field of values of A is very close to the imaginary axis. The computational cost of the general rational Krylov subspace method may be much higher than for the extended space, since a group of new shifted linear system with the same right-hand side needs to be solved at each iteration, with a different shift at each iteration. On the other hand, the extended method only requires system solves with A : if A is such that an efficient LU decomposition can be performed, then this is done once for all at the beginning of the computation and only the backward solves are required while expanding the space. The numerical experiments reported in [74] seem to ensure that the rational approximation space dimension remains usually very low, so that few systems have to be solved. The difference in computational costs per iteration is less significant if an iterative solver is used to solve the inner systems; in that case, the extended method can possibly still reuse the same preconditioner, but the computation with the iterative method still needs to be performed anew.

The rational function idea is particularly appealing when $C = C_1 C_1^\top$ has rank p larger than one. In that case, the extended Krylov subspace increases its dimension by $2p$ vectors per iteration, making the whole procedure memory consuming if convergence is not fast.

In general, memory requirements may become a serious concern when C_1 has rank much larger, say a few tens, since the approximation space dimension increases proportionally with that rank. In [75] a tangential procedure is proposed to expand the rational Krylov subspace at each iteration, in a way that only the most relevant directions are retained. More precisely, small matrices d_1, \dots, d_k are determined so that the following space is constructed

$$\text{range}([(A + \sigma_1 I)^{-1} B d_1, (A + \sigma_2 I)^{-1} B d_2, \dots, (A + \sigma_k I)^{-1} B d_k]).$$

The actual column dimension of each d_j may vary as the iteration proceeds. Numerical experiments reported in [75] show that this strategy is capable of successfully handling the presence of many columns in C_1 , and provides a performance that is largely superior to that of the block rational Krylov subspace.

The Global Krylov subspace method for the Sylvester equation was applied to the Lyapunov equation in [140], with natural simplifications due to the fact that a single space needs to be generated; numerical experiments in [140] showed better

performance than the standard block Krylov subspace methods. No numerical comparisons seem to be available for global and rational Krylov subspace methods. We also mention [121] for a generalization to the simultaneous solution of a coupled pair of Lyapunov equations, corresponding to the two Gramians of a dynamical system. This last problem was also considered in [134]: the coupled block Arnoldi method and the nonsymmetric block Lanczos recurrence were analyzed for simultaneously approximating both Gramians, so as to obtain approximations to the linear transfer function of the system; see also [136] for enhancements of the proposed approaches.

The Galerkin condition for the Lyapunov equation residual can be replaced by a Petrov-Galerkin condition; see the discussion around Proposition 4.2 for the Sylvester equation. If the constraint space is $A\mathcal{V}$, then the resulting algorithm minimizes the residual in the Frobenius norm and the problem admits the following formulation: Find $\mathbf{X}_k = V_k \mathbf{Y}_k V_k^\top$ such that \mathbf{Y}_k satisfies

$$\mathbf{Y}_k = \arg \min_{Y_k \in \mathbb{R}^{pk \times pk}} \|AV_k Y_k V_k^\top + V_k Y_k V_k^\top A^\top + C_1 C_1^\top\|_F, \quad (5.7)$$

where the columns of V_k form a basis for the approximation space \mathcal{V} . Assume once again that for the space \mathcal{V} a relation of the type $AV_k = \widehat{V}_{k+1} \underline{H}_k$ is available, with $\widehat{V}_{k+1}^\top \widehat{V}_{k+1} = I$ and \underline{H}_k of size $p(k+1) \times pk$. Then (5.7) can be rewritten in smaller dimension as the following matrix least squares problem (see, e.g., [176]):

$$\mathbf{Y}_k = \arg \min_{Y \in \mathbb{R}^{pk \times pk}} \left\| \underline{H}_k Y [I, 0] + \begin{bmatrix} I \\ 0 \end{bmatrix} Y \underline{H}_k^\top + \begin{bmatrix} \gamma_0^2 & 0 \\ 0 & 0 \end{bmatrix} \right\|_F, \quad (5.8)$$

where $C_1 = V_k \gamma_0$. This approach was explored in [135] for the standard block Krylov subspace and in [131] for the rank-one case. The projected problem entails the solution of (5.8) for which expensive (order $\mathcal{O}((pk)^4)$) procedures have been proposed [135],[131]. More recently, the minimal residual method was revisited in [176] for a generic low rank C_1 , and a more effective (order $\mathcal{O}((pk)^3)$) solver for the inner problem (5.8) was proposed.

5.2.2. ADI method. For $B = A^\top$, the ADI method of section 4.4.2 for the Sylvester equation simplifies, leading to the following recursion for the whole matrix \mathbf{X}_j with two half steps (see [255]):

$$\begin{aligned} \mathbf{X}_0 &= 0, \\ (A + s_j I) \mathbf{X}_{j-\frac{1}{2}} &= C_1 C_1^\top - \mathbf{X}_{j-1} (A^\top - s_j I_n), \\ (A + s_j I) \mathbf{X}_j &= C_1 C_1^\top - (\mathbf{X}_{j-\frac{1}{2}})^* (A^\top - s_j I_n), \quad j = 1, \dots, k. \end{aligned}$$

Here the shifts $\{s_1, s_2, \dots\}$ are complex, and are employed cyclically. If both A and C_1 are real, then the approximate solution will be real and symmetric, as long as both complex conjugates shifts are used [179]. A key idea to make the recursion amenable to large dimension matrices is to keep the solution iterate in factored form. This idea was successfully explored by Penzl in 2000 in [196] and was the basis for the software package *Lyapack* [198]; see also [35]. The resulting low-rank ADI (LR-ADI) thus determines a recurrence for the factor \mathbf{Z}_j of $\mathbf{X}_j = \mathbf{Z}_j \mathbf{Z}_j^*$ as

$$\mathbf{Z}_{j+1} = [(A^* - s_j I)(A^* + s_j I)^{-1} \mathbf{Z}_j, \sqrt{-2s_j} (A^* + s_j I)^{-1} C_1], \quad (5.9)$$

with $\mathbf{Z}_1 = \sqrt{-2s_1} (A^* + s_1 I)^{-1} C_1$: the number of columns in the factor \mathbf{Z}_j is enlarged by $\text{rank}(C_1)$ columns at each iteration. The success of the low-rank ADI approach

is related to what Penzl called the *low-rank phenomenon* in the solution \mathbf{X} , namely the already mentioned fact that the eigenvalues of \mathbf{X} tend to decay very quickly towards machine precision, so that a low rank approximation appears to be possible (see section 4.1).

The iteration matrix \mathbf{Z}_j is complex during the whole iteration, whenever some of the shifts are complex. A way to overcome this problem and to maintain real arithmetic throughout whenever A and C_1 are real is discussed in [196]; see also the more recent contribution [24].

The iteration in (5.9) requires solving systems with right-hand sides \mathbf{Z}_j and C_1 at each step j . A computational improvement to decrease the number of solves per iteration was suggested in [171] (where the LR-ADI method was called CF-ADI, in which CF stands for Cholesky Factor). In there, the columns were reordered so that only the previous iterate requires solves with a shifted matrix. The resulting recurrence is given in the following algorithm (see [171, Algorithm 2]):

Algorithm 6. Given A, C_1 and $\{s_j\}$, $j = 1, \dots, j_{\max}$

1. Set $\mathbf{z}_1 = \sqrt{-2s_1}(A + s_1I)^{-1}C_1$, $\mathbf{Z}_1 = \mathbf{z}_1$
 2. For $j = 2, \dots, j_{\max}$
 - 2.1 $\mathbf{z}_j = \frac{\sqrt{-2s_j}}{\sqrt{-2s_{j-1}}}(I - (s_{j-1} + s_j)(A + s_jI)^{-1})\mathbf{z}_{j-1}$
 - 2.2 $\mathbf{Z}_j = [\mathbf{Z}_{j-1}, \mathbf{z}_j]$
- If converged stop

At each iteration, the recurrence in Algorithm 6 thus requires system solves with a fixed number of right-hand sides corresponding to the number of columns of C_1 . As for the generation of the rational Krylov subspace, a new block of linear systems needs to be solved, as the shift varies. For a very sparse A and a small number of precomputed shifts, one could consider factorizing each of the matrices $A + s_jI$ by means of a sparse solver, and then back solve at each ADI iteration. The feasibility of this procedure is clearly problem and architecture dependent.

In [24],[34] some key relations are used to show that the residual norm can be computed efficiently. More precisely, it holds that $A\mathbf{Z}_j\mathbf{Z}_j^* + \mathbf{Z}_j\mathbf{Z}_j^*A^* + C_1C_1^* = W_jW_j^*$ where W_j is a matrix of rank p (the rank of C_1) defined as (here with real poles)

$$W_j := W_{k-1} - 2s_jQ_j, \quad W_0 = C_1$$

where $Q_1 = (A + s_1)^{-1}C_1$ and $Q_j = (I - (s_j + s_{j-1})(A + s_jI)^{-1})Q_{j-1}$, $j \geq 2$. This way,

$$\|A\mathbf{Z}_j\mathbf{Z}_j^* + \mathbf{Z}_j\mathbf{Z}_j^*A^* + C_1C_1^*\|_{\star} = \|W_jW_j^*\|_{\star} = \|W_j^*W_j\|_{\star},$$

$\star = 2, F$, where the last norm is cheap to computed and memory efficient.

Additional recent contributions were devoted to further improving the computational cost per iteration. A strategy for reducing the number of solves was proposed under the name of “modified” low-rank Smith method in [110]. The idea is to compute the singular value decomposition (SVD) of the iterate at each step and, given a dropping tolerance, to replace the iterate with its best low-rank approximation. A main ingredient is that the SVD is not recomputed from scratch; instead, it is updated after each step to include the new information and then truncated to retain only those singular values that lie above the specified tolerance. The use of the SVD exploits the fact that if $Z \approx V\Sigma U^*$ is a truncated SVD decomposition of Z , then $\mathbf{X} = ZZ^* \approx V\Sigma^2V^*$ is the truncated spectral decomposition of \mathbf{X} , so that the low

rank factor can be readily maintained. In general, the procedure reduces the number of system solves per iteration, in a way that depends on the linear independence of the new iterate columns with respect to those of previous steps. Since \mathbf{X} belongs to a rational Krylov subspace, the SVD computation determines an orthogonal basis - the columns of V associated with numerically nonzero singular values - for the generated rational space. This fact makes the truncated ADI method even closer to projection methods based on the rational Krylov space: the only difference is the way the reduced solution matrix is computed; see [76] for a formalization of this relation by means of the skeleton approximation.

A bound for the difference between the traces of the solution \mathbf{X} of the Lyapunov equation and its ADI approximation is proposed in [245], which shows that the right-hand side of the Lyapunov equation can sometimes greatly influence the eigenvalue decay rate of the solution.

Computation of the shifts. The selection of the ADI parameters and of their number has been a major topic of research for many years, since the performance of the method, in terms of number of iterations, heavily depends on them.

Let A be stable. Assuming a zero starting approximate solution, from the general ADI recurrence it follows that the error matrix associated with the ADI approximation $\mathbf{X}_k^{\text{ADI}}$ after k full iterations is given by (see also [196])

$$\mathbf{X} - \mathbf{X}_k^{\text{ADI}} = (\bar{r}_k(A)r_k(-A)^{-1})\mathbf{X}\bar{r}_k(A)^*r_k(-A)^{-*}, \quad (5.10)$$

$$r_k(z) = \prod_{i=1}^k (s_i - z), \quad \bar{r}_k(z) = \prod_{i=1}^k (\bar{s}_i - z).$$

This expression shows that for A diagonalizable, for a fixed k optimal parameters can be obtained by solving the minimax problem

$$\min_{s_1, \dots, s_k \in \mathbb{C}^-} \max_{\lambda \in \Lambda(A)} \prod_{j=1}^k \left| \frac{\lambda - s_j}{\lambda + s_j} \right|. \quad (5.11)$$

The value of k is adjusted so that the set $\{s_1, \dots, s_k\}$ is closed under conjugation in the case A is real. It is worth mentioning that it can be computationally advantageous to repeatedly apply a modest number of poles, rather than use a larger set of poles that gives a marginally faster convergence rate for the scalar rational approximation problem, if the cost of computing these poles is significant. For A having real spectrum, this minimax problem was solved by Zolotaryov; if A is also symmetric, this leads to an asymptotically optimal linear convergence rate for the approximation. The optimal parameters are then given as (see, e.g., [79])

$$s_j = \text{dn} \left[\frac{(2j-1)K}{2t}, m \right], \quad j = 1, \dots, k$$

where dn is a Jacobian elliptic function, and K is the complete elliptic integral of the first kind, of modulus m [1]. Generalizations to the case when the complex spectrum lies in certain specified complex regions Ω were discussed in [79]. However, it was only with the heuristic approach of Penzl in [196] that the computation of *suboptimal* ADI parameters became a more manageable procedure. The proposed strategy is performed as a pre-processing of the actual ADI computation: consider the Krylov

subspaces $K_{k_A}(A, c)$, $K_{k_{A^{-1}}}(A^{-1}, c)$ for some vector c , and let V, W be such their orthonormal columns span the two spaces, respectively. Let Ω_+, Ω_- be the regions containing the eigenvalues of $V^\top AV$ and of $W^\top AW$ (the Ritz values). The key idea in [196] is to replace the spectrum of A with the region $\Omega := \Omega_+ \cup \Omega_-$, and then solve the minimax problem (5.11) in Ω . The set Ω may be regarded as a reasonable approximation to the region of interest, the convex hull of the spectrum of A , and it can be more cheaply computed, especially for small $k_A, k_{A^{-1}}$; see [196] and the package [198] for more technical details. An adverse effect of this pre-processing is its computational cost: for rank-one C the cost induced by the generation of both $K_{k_A}(A, c)$, $K_{k_{A^{-1}}}(A^{-1}, c)$ for some vector c to determine the suboptimal poles is comparable to that of, e.g., the construction of the extended Krylov subspace of corresponding dimension; however, by the time good suboptimal poles are determined, the extended Krylov approach has also computed an approximate solution to the Lyapunov equation.

In spite of the significant improvements in the ADI parameter estimation, however, the method remains quite sensitive to the choice of these shifts, and performance can vary dramatically even for small changes in $k_A, k_{A^{-1}}$; see, e.g., the experiments in [224]. Adaptive strategies for pole selections such as those derived for the rational Krylov subspace in [74] are hard to obtain, since a basis for the generated space is not readily available. Nonetheless, these considerations have led to the investigation of hybrid approaches, which are described later in this section.

It was observed in [171] that the ADI method actually generates a (block) rational Krylov subspace for the given vector of shifts $\mathbf{s}_k = [s_1, \dots, s_k]$. The connection between the ADI method and the Galerkin method with the rational Krylov subspace $\mathcal{K}_k(A, C_1, \mathbf{s}_k) = \text{range}([(A+s_1I)^{-1}C_1, \dots, (A+s_kI)^{-1}C_1])$, can be made more precise, when the two methods are used with the same parameters.

THEOREM 5.6. ([76, Th.3.4]) *Assume that the field of values of A and s_j , $j = 1, \dots, k$ lie in the same half complex plane, and that C_1 has rank one. Let the columns of V form an orthonormal basis of $\mathcal{K}_k(A, C_1, \mathbf{s}_k)$, and let λ_j , $j = 1, \dots, k$ be the Ritz values of A onto $\mathcal{K}_k(A, C_1, \mathbf{s}_k)$, that is λ_j are the eigenvalues of V^*AV . Then the ADI approximation coincides with the Galerkin approximate solution \mathbf{X}_k with $K_k(A, C_1, \mathbf{s}_k)$ if and only if $s_j = \overline{\lambda_j}$, $j = 1, \dots, k$ (under a suitable index permutation for the λ_j 's).*

The condition $s_j = \overline{\lambda_j}$, $j = 1, \dots, k$ is seldom satisfied when the shifts are obtained by either an adaptive procedure or by a Penzl-style pre-processing (however, see [109] for an iterative process that approximates such a set of parameters, in the context of optimal model order reduction). We also recall that the bound of Proposition 4.4 shows that ADI cannot give much better results than the Galerkin approach with the rational space, while it is known that for poor poles ADI may give much larger residuals than in the optimal case [21]. Moreover, the lack of some form of optimality condition, e.g. orthogonality, seems to penalize the ADI idea; this problem was explored in recent works summarized in the next paragraph. Selected numerical experiments comparing the adaptive Rational Krylov subspace method and ADI can be found in [76].

Hybrid ADI methods. It was observed in [28] that “The most criticized property of the ADI iteration for solving Lyapunov equations is its demand for a set of good shift parameters to ensure fast convergence. [...] Most of the [computationally cheap parameters] are suboptimal in many cases and thus fast convergence can indeed not be guaranteed. Additionally, if the convergence is slow, the low-rank Cholesky factors

may grow without adding essential information in subsequent iteration steps.” In [28] it was thus suggested to combine the subspace projection idea with the ADI recurrence. The projection is performed onto the space spanned by the columns of the current ADI factor, the idea being motivated by the fact the ADI solution factor belongs to the rational Krylov subspace with the same shifts as ADI. The projection is performed every \hat{k} ADI iterations by computing an orthonormal basis spanning the range of the current factor, and the small size projected equation is solved by means of a Schur-type method (see section 5.1). Since the number of columns grows at each iteration, the cost of computing the orthonormal basis significantly increases. To overcome this problem, the authors suggest to truncate the obtained projected solution so that a small rank factor is retained for the next ADI iteration. More technical implementation details can be found in [28]. The idea is very reminiscent of a restarting process in the standard linear system framework, although here the truncation is performed in a different fashion. To complete this parallel with linear system solves, this procedure may be viewed as a hybrid restarted process, where a rational function (here the ADI single step) is applied to the solution before restart; see, e.g., [228] for a review of polynomial acceleration procedures of restarted methods in the linear system setting. The resulting process is called the Galerkin projection accelerated LRCF-ADI (LRCF-ADI-GP). Note that although ADI does not require that A be either positive or negative definite, the extra projection step is ensured not to break down only under the additional definiteness constraint. It is also interesting to observe that, without the truncation of the projected solution, the procedure might be mathematically equivalent to the Galerkin method in the Rational Krylov subspace obtained with the same shift parameters; a formal proof still needs to be carried out. Selected numerical experiments comparing the adaptive *tangential* Rational Krylov subspace method and projected ADI can be found in [75].

We also mention the procedure proposed in [138], where the continuous Lyapunov equation is first transformed into a discrete (Smith) Lyapunov equation with rational matrix functions as coefficient matrices, and then solved by means of the Global Krylov subspace method. This may be viewed as a preconditioning strategy.

5.2.3. Spectral, sparse format and other methods. As for the Sylvester equation, the Kronecker formulation can be used to restate the matrix equation as the following very large linear system,

$$\mathcal{A}\mathbf{x} := (I_m \otimes A + A^\top \otimes I_n)\mathbf{x} = \mathbf{c}, \quad \mathbf{x} = \text{vec}(\mathbf{X}), \quad \mathbf{c} = \text{vec}(C) \quad (5.12)$$

of size n^2 , if n is the size of A ; see, e.g., [125] for an early attempt. For A symmetric and positive definite, the convergence rate of CG applied to the Kronecker formulation is driven by the condition number $\kappa(\mathcal{A}) = \kappa(A)$, whereas the convergence rate of the Galerkin procedure directly applied to the original Lyapunov equation is driven by $\kappa(A + \lambda_{\min}I)$ (see Theorem 5.4), which can be significantly smaller than $\kappa(A)$. This analysis justifies the better performance of projection methods applied to the matrix equation. A second possibly stronger argument is given by memory requirements: the Kronecker formulation requires n^2 -length vectors. Nonetheless, it was recently shown in [183] that when solving (5.12) floating point operations can be carried out so as to lower memory storage from $\mathcal{O}(n^2)$ to $\mathcal{O}(n)$. Moreover, a standard Krylov subspace method for (5.12) can take full advantage of the structure, since most matrix-vector multiplications can be rewritten as matrix-matrix operations.

A possible way to overcome slow convergence is to choose an effective preconditioning strategy, that can drastically improve the spectral properties of the coefficient

matrix \mathcal{A} . Hochbruck and Starke used a Krylov subspace solver for the system (5.12), and they investigated SSOR and ADI iteration (with a fixed number of iterations) as operator-based preconditioners; see also [183] for some implementation aspects of preconditioning strategies. More recently, a flexible GMRES approach was proposed in [44], which allowed for a variable ADI preconditioning step. Very preliminary numerical results report promising performance of the Kronecker formulation, while taking into account the matrix structure. These approaches may have broader applications for more general matrix equations; see the discussion in section 7.2.

A rather different approach consists of using an appropriately modified version of the sign function iteration depicted in (5.2). As memory requirements are excessive in its original form for large scale problems, two major amendments have been explored (see, e.g., [17]): i) A sparsified version of A , so as to substantially reduce the computation and storage of Z_k^{-1} ; ii) For $C = C_1 C_1^\top$, a factored version of the approximation $\tilde{\mathbf{X}}$, so that only a full (low rank) factor need be iterated. The latter problem was addressed in [27], where the following coupled iteration was proposed:

$$A_0 = A, \quad B_0 = C_1, \quad A_{k+1} = \frac{1}{2}(A_k + A_k^{-1}), \quad B_{k+1} = \frac{1}{\sqrt{2}}[B_k, A_k^{-1}B_k], \quad k = 0, 1, \dots,$$

giving $\mathbf{Y} = \frac{1}{\sqrt{2}} \lim_{k \rightarrow \infty} B_k$, with $\mathbf{Y}\mathbf{Y}^\top = \mathbf{X}$. Note that the size of B_k is doubled at each iteration, therefore a rank reduction is suggested in [27]. More recent developments include new algorithms that appear to be well suited for large problems with dense and unstructured matrices [222]; the discussion in [222] in fact addresses the generalized Sylvester equation.

Item i), namely reducing the cost of dealing with the explicit inverse of large matrices, may be addressed by exploiting data-sparse matrix representations and approximate arithmetic. In [19], but also in previous related works for the algebraic Riccati equation, the \mathcal{H} -matrix format was used (see section 4.4.3). If $\text{Inv}_{\mathcal{H}}(A)$ denotes the inverse in the \mathcal{H} -matrix format, then the coupled recurrence above can be performed as

$$A_{k+1} = \frac{1}{2}(A_k + \text{Inv}_{\mathcal{H}}(A_k)), \quad B_{k+1} = \frac{1}{\sqrt{2}}[B_k, \text{Inv}_{\mathcal{H}}(A_k)B_k], \quad k = 0, 1, \dots,$$

where the sum to obtain A_{k+1} is intended in \mathcal{H} -matrix format. More implementation details can be found in [19]. According to the analysis performed in there, the error induced by the new format can be controlled while performing the rank reduction of B_{k+1} , so that the format errors do not grow unboundedly with k ; these results are in agreement with the general theory of \mathcal{H} -matrices for Riccati equations, developed in [102]. In [19], the derivation with the \mathcal{H} -matrix format is extended to the case of the generalized Lyapunov equation (see section 7). Numerical experiments show that the \mathcal{H} -format allows the sign function iteration to be employed for medium size problems ($\mathcal{O}(10000)$), for which the dense algorithm requires excessive memory allocation. Finally, an example comparing a linear multigrid solver using \mathcal{H} -format matrices with ADI is reported in [100, sec. 7.6], showing that on that specific example the multigrid approach is about ten times faster than ADI (implemented in the same framework), although both methods scale linearly with the number of multigrid levels.

We conclude this section with strategies that are more explicitly based on invariant subspaces. All considered methods assume that the maximum rank of a sufficiently accurate approximate solution is either known or given. Therefore, the context where these approaches are used is different from the one of previous methods.

The integral representation of \mathbf{X} in (4.5) and the decay of the singular values of \mathbf{X} suggest various eigenvalue-based strategies. One such method focuses on approximating the leading invariant subspace of \mathbf{X} . In [129] and [126] an Approximate Power Iteration (API) approach was proposed, which aims to approximate the dominant eigenvectors of \mathbf{X} . The method is closely related to the power iteration and the Lanczos method for computing the extremal eigenpairs of a positive definite symmetric matrix, and the authors report good convergence properties when the eigenvalues associated with the sought after eigenvectors are away from zero and well separated from the others, so that a good low-rank approximation of \mathbf{X} can be determined. The method works under the assumption that A is negative definite, as with projection methods. The API method applies the power method to \mathbf{X} , which is only known implicitly and approximately by means of products of the type $\mathbf{Y} = \mathbf{X}v$ through the solution of the auxiliary “tall” Sylvester equation

$$A\mathbf{Y} + \mathbf{Y}\Theta + q = 0, \quad (5.13)$$

where $\Theta = v^\top Av$ is a small square matrix and $q = C_1 C_1^\top v$ (see section 4.3). The numerical experiments reported in [126] on actually quite small problems, seem to imply that API is a promising method for the approximation of the leading eigenvectors of \mathbf{X} , without the computation of \mathbf{X} itself. The approach is reminiscent of the implicitly restarted Arnoldi method [169], although each iteration requires the solution of a Sylvester equation. A variant of this approach was proposed in [252] to overcome misconvergence caused by the omission of the term $\mathbf{X}(I - vv^\top)A^\top v$ in the equation (5.13). Motivated by [126], an algorithm combining the power method and (implicitly) the ADI iteration was proposed in [188]; see [187] for a more thorough presentation of these approaches.

With the same aim of approximating the leading invariant subspace of \mathbf{X} of given dimension, the procedure explored in [107] performs a refined numerical approximation by repeatedly integrating the dynamical system associated with the Lyapunov equation, as the basis for an orthogonal power iteration.

A somewhat related approach was proposed in [229], and it exploits the popular *proper orthogonal decomposition* (POD) approach employed in reduced order modeling of large-scale dynamical systems [26]. The idea is to collect a sample of m solutions to the associated linear differential equation with different starting data, and for a chosen k , form a rank- k approximate Lyapunov solution by using the eigenpairs of the sample, obtained by collecting the approximate solutions of a sequence of linear time-dependent differential equations with different starting solutions (the so-called method of snapshots). The approach relies on the integral representation of the Lyapunov solution, and according to the author, it is particularly appropriate for infinite dimensional problems.

Finally, a novel very different approach was recently proposed by Vandereycken and Vandewalle in [250] for A symmetric and positive definite: the method finds a low-rank approximation to \mathbf{X} by minimizing the function

$$f : \mathcal{M}_k \rightarrow \mathbb{R}, \quad X \mapsto \text{trace}(XAX) - \text{trace}(XC)$$

on the manifold of symmetric and positive semidefinite matrices of rank k in $\mathbb{R}^{n \times n}$, namely

$$\min_{X \in \mathcal{M}_k} f(X).$$

Note that the minimization of f corresponds to the minimization of the error in the energy norm, having used the Kronecker formulation of the Lyapunov equation. By using the smoothness of \mathcal{M}_k the problem is solved within a Riemann optimization framework, which allows one to embed the rank constraint in the space, and solve an unconstrained minimization problem by means of a Riemann trust-region method, a second-order model based on the Hessian [2]. At convergence of the minimization process, if the current solution rank is not sufficiently accurate, the process is restarted basically from scratch. As a result, the method may be appealing when the optimal rank is approximately known a-priori, otherwise the approach may not be competitive with respect to other strategies discussed so far.

6. The Stein and discrete Lyapunov equations. The Stein and the discrete Sylvester equations are the discrete-time counter part of the (continuous-time) equations discussed in the previous sections, and they naturally stem from a discrete-time system; see (3.2) and, e.g., [4, sec.4.3]. Other relevant applications include for instance statistics [151],[150], probability [10], spectral analysis [132]; these equations are also a computational tool in the design of control systems [117], or in the coprime matrix fraction description of linear systems [271].

The Stein equation may be written as

$$\mathbf{X} + \mathbf{A}\mathbf{X}\mathbf{B} = \mathbf{C}, \quad (6.1)$$

where it is assumed that the eigenvalues of A and B are contained in the open unit disk. The discrete-time Lyapunov equation is obtained by choosing $B = -A^\top$, in which case, if C is symmetric and if a solution \mathbf{X} exists, then \mathbf{X} has to be symmetric. In the context of inertia theory, for $C \succeq 0$ the discrete-time Lyapunov equation allows one to analyze the proximity of $\text{spec}(A)$ to the unit circle and the proximity of $\text{spec}(X)$ to the imaginary axis; see, [165, sec.13.2], and also, e.g., [264],[170] for more specialized results.

Under the condition that $\lambda_i(A)\lambda_j(B) \neq -1$ for all i, j , the solution \mathbf{X} exists and is unique for any C (see, e.g., [164]), and this is highlighted by the Kronecker form of (6.1), given as $(I + B^\top \otimes A)\mathbf{x} = c$, where $\mathbf{x} = \text{vec}(\mathbf{X})$ and $c = \text{vec}(C)$. Necessary and sufficient conditions for the existence and uniqueness of the solution \mathbf{X} were obtained in [263] as a generalization of the property (4.2) for the Sylvester equation. Inertia and other transformation-based results for $B = -A^\top$ can be derived in a natural manner from those on the Lyapunov equation; see, e.g., [219],[220]. We also refer the reader to [163] for a solution expressed in terms of the companion form of the given matrices, and to [39] for related computational considerations.

To numerically solve the equation for, say, B nonsingular, one could think of working with $\mathbf{X}B^{-1} + \mathbf{A}\mathbf{X} = \mathbf{C}B^{-1}$ which is a standard Sylvester equation, and then adopt one of the solution methods from previous sections. In fact, (6.1) is nothing but a generalized Sylvester equation as in (1.1) with special choices of the first two coefficient matrices. For large B , the matrix B^{-1} should not be formed explicitly, but its action used within iterative methods.

Forming B^{-1} explicitly is not recommended also in the small size case, whenever B is ill-conditioned. Alternative transformations that bring the discrete equation to standard form are given by (here for $B = -A^\top$, see [202])

$$\tilde{\mathbf{A}}\tilde{\mathbf{X}} + \tilde{\mathbf{X}}\tilde{\mathbf{A}}^\top = \mathbf{C}, \quad \text{with } \tilde{\mathbf{A}} = (A - I)^{-1}(A + I), \quad \mathbf{X} = \frac{1}{2}(\tilde{\mathbf{A}} - I)^\top \tilde{\mathbf{X}}(\tilde{\mathbf{A}} - I),$$

and ([12], [199], [147])

$$\tilde{A}\mathbf{X} + \mathbf{X}\tilde{A}^\top = \tilde{C}, \quad \text{with } \tilde{A} = (A - I)(A + I)^{-1}, \tilde{C} = 2(A^\top + I)^{-1}C(A + I)^{-1},$$

where it is assumed that the inversions are well defined. In general, however, the same stability considerations as for methods using A^{-1} apply.

All these difficulties encourage solving the discrete equations (6.1) directly. A Schur-form type method for small size coefficient matrices that directly deals with (6.1) can be found, e.g., in [13], while a generalization of the ‘‘continuous-time’’ Hammarling method was proposed by Hammarling himself in [114].

In [251], Varga established a rank-two updating formula for the Cholesky factors in Hammarling’s algorithm for solving the real, non-negative definite Stein equation. As a generalization of his algorithm for the Lyapunov equation, a block variant of Hammarling’s method for the discrete-time Lyapunov equation is suggested in [156].

In spite of the large similarity with the standard equation, directly attacking (6.1) is an interesting problem in its own, especially for A and B of large dimensions and with either of the two matrices singular. For a low rank \tilde{C} , projection methods are applicable to solve (6.1), and an approximate solution $\tilde{\mathbf{X}} = V_k \mathbf{Y} W_k^\top$ can be determined, where the columns of V_k and W_k span approximation spaces associated with A and B^\top , respectively. For instance, a global Krylov subspace approach was proposed in [137, sec. 5],[143], and its implementation is a natural modification of that used for the standard Sylvester equation. Similar derivations can be obtained for other Krylov-based methods.

The discrete-time Lyapunov equation motivated the development of the Smith method [231], which is at the basis of the modern ADI iteration for the Lyapunov equation. For A d -stable (i.e., with eigenvalues inside the unit circle), the unique solution to (6.1) with $B = -A^\top$ can be written as $\mathbf{X} = \sum_{j=0}^{\infty} A^j C (A^j)^\top$, and it is symmetric and positive semi-definite, if C is. The (convergent) Smith iteration is defined as

$$\mathbf{X}_0 = 0, \quad \mathbf{X}_{k+1} = C + A\mathbf{X}_k A^\top,$$

with a closed form given by $\mathbf{X}_k = \sum_{j=1}^k A^{j-1} C (A^{j-1})^\top$. Faster - quadratic - convergence can be achieved with the *squared* Smith method, which becomes of interest in the large scale case precisely for C of small rank [196]. The iteration is generically given as

$$\mathbf{X} = A^{2^{k+1}} \mathbf{X} (A^{2^{k+1}})^\top + \sum_{i=0}^{2^{k+1}-1} A^i C (A^i)^\top, \quad \mathbf{X} = \lim_{k \rightarrow \infty} \sum_{i=0}^{2^{k+1}-1} A^i C (A^i)^\top.$$

The resulting recursion is given by $H_{k+1} = H_k + A_k H_k A_k^\top$, $H_0 = C$, where $A_{k+1} = A_k^2$, so that $C_k \rightarrow \mathbf{X}$ as $k \rightarrow \infty$. By exploiting the low rank of $C = C_0 C_0^\top$, $H_{k+1} = C_{k+1} C_{k+1}^\top$ with $C_{k+1} = [C_k, A_k C_k]$. Therefore, the number of columns of C_{k+1} doubles at each iteration, and C_{k+1} is contained in a block Krylov subspace generated by A and C_0 . Recent advances to make this recurrence more effective both in terms of computational costs and memory requirements include compressions, truncations and restarts, with a tricky use of the underlying Krylov subspace [173],[214],[31]. In these references, estimates for the residual and error norms are also derived. Finally, we point out an ADI acceleration strategy in [214] (for $B = -A^\top$) and in [31], which significantly improves the convergence speed. In fact, a major breakthrough for the

Smith method consisted in combining its recurrence with the ADI idea, as developed in [196].

All these approaches rely on the fact that often the solution \mathbf{X} has (numerical) rank much lower than n ; indeed, in [31] it is shown for the Stein equation that if the eigenvalues of A and B lie inside the open unit disk and C has rank p ,

$$\frac{\sigma_{kp+1}(\mathbf{X})}{\sigma_1(\mathbf{X})} \leq \|A^k\| \|B^k\|,$$

indicating that indeed the solution rank may be small if the powers of A and B decrease rapidly in norm. In [214] the following estimate was derived for $B = -A^\top$ and $\|A\| < 1$:

$$\frac{\sigma_{kp+1}(\mathbf{X})}{\sigma_1(\mathbf{X})} \leq \frac{\|A\|^{2k}}{1 - \|A\|^2}.$$

In general, a computational comparison of various variants of the approaches based on the Smith iteration is still lacking, though it would be highly desirable.

A related matrix equation is the \top -Stein equation, given by $\mathbf{X} = A\mathbf{X}^\top B + C$, whose solvability conditions have been recently analyzed in [175]. More generally, a broader class of matrix equations can be written as $\mathbf{X} = Af(\mathbf{X})B + C$, where $f(X) = X^\top$, $f(X) = \bar{X}$ or $f(X) = X^*$, whose analysis and numerical solution can be recast in terms of the Stein matrix equation [272]. This and more general forms of linear equations are discussed in the next section.

7. Generalized linear equations.

7.1. The Generalized Sylvester and Lyapunov equations. The term *generalized* refers to a very wide class of equations, which includes systems of matrix equations, bilinear equations and problems where the coefficient matrices are rectangular. We start with the most common form of generalized Sylvester equation, namely

$$AXD + EXB = C, \tag{7.1}$$

which differs from (4.1) for the occurrence of coefficient matrices on both sides of the unknown solution \mathbf{X} .

If D and E are both nonsingular, left multiplication by E^{-1} and right multiplication by D^{-1} lead to a standard Sylvester equation, with the same solution matrix \mathbf{X} . In case either E or D are ill-conditioned, such a transformation may lead to severe instabilities. This problem is common to other *generalized* equations we will encounter later in this section, and it justifies the development of solution methods that stick to the original form (7.1). The case of singular D and E , especially for $D = E^\top$ and $B = A^\top$ has an important role in the solution of differential-algebraic equations and descriptor systems [162]. The solution of (7.1) for E and D singular requires knowledge of the spectral projectors onto the right and left deflating subspaces of the *stable* pencils $\lambda E - A$ and $\lambda D - B$, associated with the finite eigenvalues along the right and left deflating subspaces associated with the eigenvalue at infinity. In such a setting, the right-hand side matrix is also projected onto the corresponding deflating subspaces, and the equation takes the name of *projected* Sylvester equation. The numerical treatment of this matrix equation necessitates of ad-hoc procedures, that appropriately and stably take into account the Weierstrass canonical form of the

pencils $\lambda E - A$, $\lambda D - B$, from which the spectral projectors can be derived; we refer the reader to, e.g., [237],[177] and their references for further details on projected Sylvester equations.

The following result ensures the existence of a unique solution \mathbf{X} to (7.1).

THEOREM 7.1. ([254]) *The matrix equation $A\mathbf{X}D + E\mathbf{X}B = C$ has a unique solution if and only if*

- (i) *The pairs (A, E) and $(D, -B)$ are regular pencils;*
- (ii) *The spectra of (A, E) and $(B, -D)$ are disjoint¹³.*

Under the hypotheses of Theorem 7.1, uniqueness is thus still ensured if one of the matrices A, B, D or E is singular, as long as the corresponding pencil is nonsingular.

A natural extension of the Bartels-Stewart method can be implemented for numerically solving (7.1) when dimensions are small, and this was discussed in [91], [92], [194], where the starting point is a QZ decomposition of the pencils (A, E) and (B, D) , followed by the solution of a sequence of small (1-by-1 or 2-by-2) generalized Sylvester equations, which is performed by using their Kronecker form. For C positive semidefinite and (A, E) stable, in [194] a generalization of the Hammarling method is also proposed. The algorithm developed in [91], [92] is also able to treat the case in which some specifically selected coefficient matrices are singular.

The large scale setting does not significantly differ from previous cases, as long as E, D are not too ill-conditioned. The problem can be recast as a standard Sylvester equation in $E^{-1}A$ and BD^{-1} . In the case of rational Krylov subspace and ADI methods, shifted systems can be solved with the coefficient matrix $(E^{-1}A + sI) = E^{-1}(A + sE)$, and analogously for systems with BD^{-1} . In the case of ill-conditioned E, D , one could consider using a specifically selected $\alpha \in \mathbb{R}$ (or $\alpha \in \mathbb{C}$) such that the two matrices $E + \alpha A$ and $D - \alpha B$ are better conditioned and the solution uniqueness is ensured, and rewrite (7.1) as the following equivalent generalized Sylvester matrix equation, $A\mathbf{X}(D - \alpha B) + (E + \alpha A)\mathbf{X}B = C$.

We mention the specific application of global Krylov subspace methods (see section 4.4), which are obtained by using the mapping $\mathcal{M}(X) = AXD + EXB$, therefore they can be applied in general to the equation $\sum_{i=1}^q A_i X B_i = C$, as done in [46]; note that this kind of approach can only be applied to medium size problems, as the matrix formulation involves dense matrices. We recall once again that there is a tight relation between global methods and the Kronecker form, which provides a good ground for the theoretical understanding of global methods.

A unique solution to the generalized Lyapunov equation

$$A\mathbf{X}E^\top + E\mathbf{X}A^\top = C, \quad (7.2)$$

is ensured if and only if Theorem 7.1 applies, that is all eigenvalues of the pencil (A, E) are finite, and they do not have pairwise zero sum. As a consequence, a unique solution is only obtained if one of the matrices A, E is nonsingular. In this case one can recast equation (3.3) as a standard Lyapunov equation.

To avoid stability problems caused by a possibly ill-conditioned E or A , it is usually preferred to work with E and A implicitly. This is realized by performing a simultaneous Schur decomposition of E and A , $E = QSZ^*$, $A = QTZ^*$, with S and

¹³Here the notion of disjoint spectra ([254, formula (7)]) should be intended keeping in mind the definition of “spectral set” in generalized eigenvalue problems, as defined for instance in [236, Definition VI.1.1].

T (complex) upper triangular [184]. Plugging in this transformation, (3.3) becomes $QTZ^*XZS^*Q^* + QSZ^*XZ^*T^*Q^* = C$, that is

$$T\widehat{\mathbf{X}}S^* + S\widehat{\mathbf{X}}T^* = Q^*CQ, \quad \widehat{\mathbf{X}} = Z^*XZ.$$

The elements of $\widehat{\mathbf{X}}$ can then be obtained by exploiting the structure of T and S [113].

A different approach adapts the matrix sign function iteration in (5.3) to this more general context, and it is shown in [27] that it is applicable under the hypothesis that the Lyapunov equation is stable. In the case of C in factorized form in (3.3), a recurrence is proposed in [27] to generate an approximation to the Cholesky-type factor of the resulting semidefinite solution \mathbf{X} . Comparisons in terms of memory requirements and floating point operations with respect to the generalized Hammarling method (see [194]) are also reported in [27]. We also refer the reader to [194] for some estimates of the separation¹⁴ and the condition number of the operator associated with (3.3), which is important to assess the accuracy of the computed solution.

7.2. Bilinear, constrained, and other linear equations. Other generalizations of the Sylvester equation have attracted the attention of many researchers. In some cases the standard procedure for their solution consists in solving a (sequence of) related standard Sylvester equation(s), so that the computational core is the numerical solution of the latter by means of some of the procedures discussed in previous sections. We thus list here some of the possible generalizations more often encountered and employed in real applications. We start by considering the case when the two coefficient matrices can be rectangular. This gives the following equation:

$$\mathbf{A}\mathbf{X} + \mathbf{Y}\mathbf{B} = C, \tag{7.3}$$

where \mathbf{X} , \mathbf{Y} are both unknown, and A , B and C are all rectangular matrices of conforming dimensions. Equations of this type arise in control theory, for instance in output regulation with internal stability, where the matrices are in fact polynomial matrices (see, e.g., [262] and references therein). The following theorem is a first result on the existence and uniqueness of the pair \mathbf{X} , \mathbf{Y} , and is reported as originally stated in [209]; see also more recent advanced developments in [81].

THEOREM 7.2. ([209]) *The necessary and sufficient condition that the equation $\mathbf{A}\mathbf{X} - \mathbf{Y}\mathbf{B} = C$, where A, B , and C are $m \times r$, $s \times n$, and $m \times n$ matrices respectively with elements in a field \mathcal{F} , have a solution \mathbf{X}, \mathbf{Y} of order $r \times n$ and $m \times s$ respectively and with elements in \mathcal{F} is that the matrices*

$$\begin{bmatrix} A & C \\ 0 & B \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} A & 0 \\ 0 & B \end{bmatrix}$$

be equivalent.

The matrix equivalence in the theorem can be explicitly obtained as

$$\begin{bmatrix} I & \mathbf{Y} \\ 0 & I \end{bmatrix} \begin{bmatrix} A & C \\ 0 & B \end{bmatrix} \begin{bmatrix} I & -\mathbf{X} \\ 0 & I \end{bmatrix} = \begin{bmatrix} A & 0 \\ 0 & B \end{bmatrix},$$

where \mathbf{Y} and \mathbf{X} are the solution matrices to the matrix equation (7.3). Solvability conditions of this equation can be derived from [236, Chapter VI]; the result is stated next.

¹⁴Defined as $\text{sep}_p(A, E) = \min_{\|X\|_p=1} \|A^\top XE + E^\top XA\|_p$, with $p = 2, F$.

PROPOSITION 7.3. ([236]) *Assume A and B are nonsingular. Then the problem $\mathbf{A}\mathbf{X} + \mathbf{Y}B = C$ has a unique solution when the spectra of A and B are disjoint.*

These are extensions of the analogous result for the standard Sylvester equation, see (4.2) and [209], [130]. Note that by setting for instance $\mathbf{U} = [\mathbf{X}; \mathbf{Y}]$, equation (7.3) can be rewritten as a

$$\mathbf{A}\mathbf{U} \begin{bmatrix} I_n \\ 0 \end{bmatrix} + \mathbf{U} \begin{bmatrix} 0 \\ B \end{bmatrix} = C$$

in the single unknown matrix \mathbf{U} [155].

The two-sided version of (7.3) is given by

$$\mathbf{A}\mathbf{X}D + E\mathbf{Y}B = C,$$

and this is an example of more complex bilinear equations with several left-hand side terms considered in the literature; see, e.g., [271] and references therein.

A typical generalization is given by the following bilinear equation:

$$\mathbf{A}\mathbf{X}D + E\mathbf{X}B = C\mathbf{Y} + F, \quad (7.4)$$

where the pair (\mathbf{X}, \mathbf{Y}) is to be determined, and \mathbf{X} occurs in two different terms. Theoretical aspects are collected in [265], and also in [266], where closed forms for (\mathbf{X}, \mathbf{Y}) . In [270] general parametric expressions for the solution matrices \mathbf{X} and \mathbf{Y} are also obtained, under the hypothesis that D is full rank and F is the zero matrix.

The main objective in the aforementioned papers is in fact the solution of *systems* of bilinear matrix equations:

$$\begin{cases} A_1\mathbf{X} + \mathbf{Y}B_1 = C_1 \\ A_2\mathbf{X} + \mathbf{Y}B_2 = C_2; \end{cases} \quad (7.5)$$

see, e.g., [146],[77], for which a recent perturbation analysis can be found in [178]. These systems can arise, for instance, in the numerical treatment of systems of stochastic partial differential equations, giving rise to large and sparse coefficient matrices; see, e.g., (3.13) and [83]. The system (7.5) is an important step in deflating subspace computations for pencils [144],[236, Ch.VI],[254]. Indeed, the system can be formulated in terms of a transformation $\mathcal{P}^{-1}(\mathcal{M} - \lambda\mathcal{N})\mathcal{Q}$ onto a block diagonal form, of the matrix pencil

$$\mathcal{M} - \lambda\mathcal{N} = \begin{bmatrix} A_1 & -C_1 \\ 0 & -B_1 \end{bmatrix} - \lambda \begin{bmatrix} A_2 & -C_2 \\ 0 & -B_2 \end{bmatrix}.$$

The pair (\mathbf{X}, \mathbf{Y}) is sought such that

$$\mathcal{P}^{-1}(\mathcal{M} - \lambda\mathcal{N})\mathcal{Q} := \begin{bmatrix} I & -\mathbf{Y} \\ 0 & I \end{bmatrix} (\mathcal{M} - \lambda\mathcal{N}) \begin{bmatrix} I & \mathbf{X} \\ 0 & I \end{bmatrix} = \begin{bmatrix} A_1 & 0 \\ 0 & -B_1 \end{bmatrix} - \lambda \begin{bmatrix} A_2 & 0 \\ 0 & -B_2 \end{bmatrix}.$$

Setting

$$\mathcal{L} : (\mathbf{X}, \mathbf{Y}) \mapsto (A_1\mathbf{X} + \mathbf{Y}B_1, A_2\mathbf{X} + \mathbf{Y}B_2),$$

the problem of solvability of (7.5) corresponds to that of determining when the operator \mathcal{L} is nonsingular. It turns out that for (A_i, B_i) , $i = 1, 2$ regular pairs, \mathcal{L} is

nonsingular if and only if the spectra of (A_1, B_1) and (A_2, B_2) are disjoint [236, Theorem 1.11]; this result also leads to the one in Proposition 7.3. The operator \mathcal{L} is a generalization of the corresponding operator for the Sylvester equation; its sensitivity, in terms of distance to singularity can be defined analogously to the sep operator; see, e.g., [146],[236, section VI.2.4].

From a numerical standpoint, the most reliable approach for small scale computations was proposed in [146] and further developed in [145], and it is based on the stable generalized Schur method, which applies the QZ algorithm to the pairs of coefficient matrices; a perturbation analysis is also included. Few other alternatives have been explored so far, that go beyond a cleverly implemented Kronecker formulation. The idea suggested in [77] amounts to “expanding” the two equations into a single one of larger size, whose solution contains both \mathbf{X} and \mathbf{Y} , but requires the Jordan decomposition of some of the coefficient matrices. A similar framework is used in [69] where more than two unknown matrices are allowed, and an approximate solution is obtained by means of a least squares approach. It is not clear how any of these procedures can be adapted to the large scale setting.

The number of linear matrix equations and unknown matrices can in fact be quite large, as discussed for instance in [45]. Necessary and sufficient conditions for the resulting systems to have a solution pair are studied in [258]. Computationally speaking, this general case has only been treated so far by using the Kronecker formulation, so that only very small problems have been tackled; but see [271], where the problem of solving the set of matrix equations is recast as an optimization problem.

A special class of *nonlinear* problems is given by the following Sylvester-Observer equation, which stems from the problem of determining a *reduced* order observer model [70], [180]. Find matrices \mathbf{X} , \mathbf{Y} and \mathbf{Z} such that

$$\mathbf{X}A - \mathbf{Y}\mathbf{X} = \mathbf{Z}C, \quad \begin{bmatrix} \mathbf{X} \\ C \end{bmatrix} \text{ invertible}, \quad (7.6)$$

where A and C are known matrices with C having few rows. A solution to (7.6) exists for any choice of spectrum of \mathbf{Y} , therefore this spectrum can be predetermined; a choice that makes \mathbf{Y} a stable matrix also ensures convergence of the reduced order observer; we refer the reader to [180] for more details on these aspects. A possible way to address the solution of (7.6) is to choose \mathbf{Y} and \mathbf{Z} arbitrarily and then solve for \mathbf{X} the resulting Sylvester equation. Early approaches in this direction did not lead to a numerically stable method. For small size matrices, the reduction to Hessenberg form proposed by Van Dooren in [70] is still one of the most effective methods for solving (7.6). The algorithm is based on a reduction to “staircase form” of the pair (A, C) , and on the determination of the solution \mathbf{X} with a particular structure in a different coordinate system. We also refer the reader to [249] for a more detailed survey on methods for dense matrices. More recently, other approaches have been proposed: for instance, a block generalization of the method in [70] was proposed in [54]; moreover, in [55] the authors proposed a block algorithm for determining a full rank solution, and it seems to be most appropriate for large-scale problems with sparse A . In this latter setting, a successful early method was proposed in [63]. The approach first assumes that $\mathbf{Z}C$ is rank-one and exploits the resemblance between the observer-Sylvester equation and the Arnoldi relation (4.17). As a by-product of the method, the authors in [63] also derive an algorithm for solving the partial-pole-assignment problem for large and sparse A , which is generalized in [65] to higher rank of $\mathbf{Z}C$. The authors in [52] propose a new strategy for a-priori choosing the eigenvalues of \mathbf{Y}

that makes the algorithm in [63] more efficient. From a control theory point of view, the possibility to determine a reduced order model is also important in the derivation of stable closed-loop systems, giving rise to a well exercised eigenvalue assignment problem. We refer the reader to, e.g., [66], for a brief survey on this and other related problems.

Within the Sylvester-observer equation, we can formulate the problem in a slightly different manner, namely by means of a constraint (see, e.g., [239],[246],[185]), and it can be stated as follows (see, e.g., [11]): *Given $A \in \mathbb{R}^{n \times n}$ stable, $B \in \mathbb{R}^{n \times p}$, $C \in \mathbb{R}^{m \times n}$ and $F \in \mathbb{R}^{(n-m) \times (n-m)}$, find $\mathbf{Y} \in \mathbb{R}^{(n-m) \times m}$ and $\mathbf{X} \in \mathbb{R}^{(n-m) \times n}$ such that*

$$\mathbf{X}A - F\mathbf{X} = \mathbf{Y}C, \quad \text{with } \mathbf{X}B = 0, \quad (7.7)$$

so that $[\mathbf{X}; C]$ is full rank.

The problem may be viewed as a homogeneous system of linear matrix equations, a generalization of (7.5), with two terms in \mathbf{X} as in (7.4), and $C_1 = 0$, $C_2 = 0$; however, there is no need to expand it by means of the Kronecker product. In [11] the authors provide necessary and sufficient conditions for a solution to (7.7) to exist, and propose an algorithm for its computation in the small-scale case. The main ingredients are a QR factorization and the solution of a standard Sylvester equation. In [93] a modification of this method was presented, to handle the case of almost singular Sylvester equations. The large scale case has been recently addressed in [218] by generalizing the method in [11]. We also point to [56] for an approach that handles a ‘‘regional pole-placement constraint’’ on F in (7.7) for a descriptor system, and to [62] and its references for further theoretical properties.

Going back to a single unknown matrix, other ‘‘generalized’’ Lyapunov equations more in the spirit of (1.2) include extra linear terms:

$$A\mathbf{X} + \mathbf{X}A^* + N\mathbf{X}N^* + C_1C_1^* = 0, \quad (7.8)$$

and they stem, for instance, from complex dynamical systems like the one in (3.7). We refer the reader to [105] for sufficient conditions on the existence of the controllability and observability Gramians; more complex forms involve more structured matrices N , see, e.g., [22],[58]. In fact, more terms of the type $N_j\mathbf{X}N_j^*$, $j = 1, 2, \dots$ could arise to fulfill more general model requests. Polynomial and infinite dimensional systems are also of interest, see, e.g., [267] and [60], respectively, and their references. In addition to a robust Kronecker-form based iteration reviewed in [61, sec.3.1-4], Damm in [61] proposed a regular splitting for the numerical solution of (7.8), yielding the following iterative scheme:

$$A\mathbf{X}_{k+1} + \mathbf{X}_{k+1}A^* = -N\mathbf{X}_kN^* - C_1C_1^*, \quad \mathbf{X}_0 = 0,$$

which entails the solution of a sequence of standard Lyapunov equations. Convergence to \mathbf{X} is obtained if the spectrum of A is sufficiently away from the imaginary axis. In [61, sec. 4] the generalized case of the Lyapunov operator is also treated. In the recent article [30] a thorough discussion and contextualization of the algebraic problem in stochastic model order reduction can be found. In [22], various methods for the Lyapunov equation, such as ADI and projection techniques, are adapted to the setting of (7.8), including sparse format approaches for the Kronecker formulation; reported experimental results on large problems seem to favor this latter approach, with the caveat of tuning the sparsity and accuracy parameters, as described in section 5.2.3.

An approach that may be appropriate for large scale problems is implicitly suggested in [8]. In the context of model order reduction, the following approximation space is introduced:

$$\text{range}(V) = \text{span} \left\{ \bigcup_{k=1}^r \text{range}\{V^{(k)}\} \right\}, \quad (7.9)$$

with $\text{range}(V^{(1)}) := K_q(A^{-1}, A^{-1}C_1)$ and

$$\text{range}(V^{(k)}) := K_q(A^{-1}, A^{-1}NV^{(k-1)}), k = 2, \dots, r.$$

Using a Galerkin approximation onto $\text{range}(V)$, the equation (7.8) can be reduced and solved with a direct procedure; a possible implementation of this idea was recently proposed in [29]. Another approach for solving multilinear systems in Kronecker form was analyzed in [161], for which a tensor-based form for the approximate solution is considered. Such a strategy is well suited in the approximation of *parameterized* linear systems, as they arise for instance in certain discretization strategies for the numerical solution of stochastic partial differential equations [6]. Data sparse methods associated with the Kronecker formulation may provide a possible successful venue for attacking the general linear multi-term matrix equation (1.2); to the best of our knowledge, no attempts have been made in this direction so far for really large problems.

7.3. Sylvester-like and Lyapunov-like equations. Sylvester and Lyapunov-like linear matrix equations of the form ([47], [40])

$$B\mathbf{X} + f(\mathbf{X})A = C, \quad A^*\mathbf{X} + f(\mathbf{X})A = C, \quad B, A, \mathbf{X} \in \mathbb{C}^{m \times n}, \quad (7.10)$$

with $f(X) = \bar{X}$, $f(X) = X^\top$, $f(X) = X^*$, or their “discrete-time” variants (see section 6) are less common; but see, for instance, [160] for an occurrence in structured eigenvalue computation. The homogeneous case ($C = 0$) has been recently analyzed in [243], where a complete description of the solution in terms of the Kronecker canonical form of $A + \lambda f(B)$ is derived, whenever information on this latter pencil is available. These equations have attracted increasing interest in the past few years, with recent contributions on the necessary and sufficient conditions for the solvability of this equation, for any right-hand side matrix C [133]; a different proof of this result that also induces a numerical method is proposed in [253]. As a sample of this type of result, in [51, Lemma 5.10] for $f(X) = X^\top$, it is proved that a unique solution \mathbf{X} exists if and only if the pencil $A - \lambda B^\top$ is regular and if its spectrum is \top -reciprocal free¹⁵, with possibly the only exception of the unit eigenvalue, which should be simple.

Interestingly, it was recently shown that for A and B nonsingular, the problem $A\mathbf{X} + \mathbf{X}^\top B = C$ can be recast as a standard Sylvester equation. The following result is proved in [71].

PROPOSITION 7.4. *Assume that A and B are nonsingular. If \mathbf{X} is a solution to the matrix equation $A\mathbf{X} + \mathbf{X}^\top B = C$, then \mathbf{X} is also a solution to the Sylvester matrix equation*

$$(B^{-\top}A)\mathbf{X} - \mathbf{X}(A^{-\top}B) = B^{-\top}C - B^{-\top}CA^{-\top}B. \quad (7.11)$$

The reverse also holds if (7.11) admits a unique solution \mathbf{X} .

¹⁵A set of complex numbers $\{\lambda_1, \dots, \lambda_k\}$ is \star -reciprocal free if $\lambda_i \neq 1/\lambda_j^\star$ for any $1 \leq i, j \leq k$. Typically, $\star = \top$ or $\star = *$, so that λ_j^\star is λ_j or $\bar{\lambda}_j$, resp.

Under the given hypotheses, this result allows one to solve the \top -Sylvester equation by means of procedures for the standard Sylvester equation, with the caveat of maintaining good stability properties of the problem.

Going back to the original formulation in (7.10), in [242] an algorithm that relies on the generalized Schur decomposition of the pair $(A, f(B))$ (via the QZ algorithm) is proposed to determine \mathbf{X} for small A and B . For $f(X) = X^\top$ this can be briefly summarized as follows:

- 1) Decompose $A = URV$ and $B^\top = USV$, with U, V unitary and R, S upper triangular;
- 2) Compute $E = VCV^\top$;
- 3) Solve $S^\top W + W^\top R = E$ element-by-element;
- 4) Form $X = \bar{U}W\bar{V}$.

The solution of the equation in step 3) is also treated in detail in [242].

Recent developments have considered the case where both A and B have large dimensions. In particular, in [71] projection methods are derived to solve the \top -Sylvester equation for the case when A and B are nonsingular. They generate right and left approximation spaces \mathcal{V}_m and \mathcal{W}_m , respectively, satisfying $B^\top \mathcal{V}_m = \mathcal{W}_m$, so that a suitable Petrov-Galerkin condition can be imposed. The reduction yields a small Sylvester-like equation of the same form, which can be solved with the Schur decomposition strategy above. The two approximation spaces can be chosen as any of the Krylov-based spaces described in previous sections; we refer the reader to [71] for algorithmic details.

In [268] a closed-form solution to the equation for $f(X) = \bar{X}$ is considered, together with the set of all possible solutions for (7.4) and for the bilinear problem $A\bar{\mathbf{X}} + B\mathbf{Y} = \mathbf{X}F$. Chiang, Duan, Feng, Wu and their collaborators have thoroughly investigated these formulations and their role in control applications.

A particular extension of this class of problems is given by polynomial equations. Consider a polynomial matrix $R(\xi) = R_0 + R_1\xi + \dots + R_\ell\xi^\ell$ in the unknown ξ , where R_i are constant square matrices and such that $\det(R(\xi))$ is not identically zero, and let Z be a square polynomial matrix satisfying $Z(\xi) = Z(-\xi)^\top$. The equation

$$R(-\xi)^\top \mathbf{X}(\xi) + \mathbf{X}(-\xi)^\top R(\xi) = Z(\xi),$$

in the square polynomial matrix \mathbf{X} is called the polynomial Sylvester equation. This special equation plays a role in the computation of integrals of quadratic functions of the variables of a system and their derivatives (see, e.g., [240]), and in the stability theory for high-order differential equations. In [193] the authors focus on the case when the right-hand side has the form $Z(\xi) = Q(-\xi)^\top \Sigma Q(\xi)$, where Q is a real rectangular polynomial matrix in ξ such that QR^{-1} is a matrix of strictly proper rational functions and Σ is a diagonal, signature matrix. An iterative solution method inspired by the method of Faddeev for the computation of the matrix resolvents is also described; we refer the reader to [115] for a detailed derivation of the Faddeev sequence in connection with the solution of Lyapunov and Sylvester equations. More general equations include polynomial Diophantine matrix equations in the form $D(\xi)\mathbf{X}(\xi) + N(\xi)\mathbf{Y}(\xi) = F(\xi)$; in [117] closed form solutions are presented, which could be used to numerically solve small size equations. In the large scale setting, this problem is computationally unsolved.

Finally, special attention should be paid to the homogeneous version of the Sylvester-like equation previously discussed with $B = A$,

$$A\mathbf{X}^\top + \mathbf{X}A = 0. \tag{7.12}$$

For each fixed complex matrix A , the solution space to this latter equation is a Lie algebra equipped with Lie bracket $[X, Y] := XY - YX$. We refer the reader to the recent articles [241], [90] and their references for more details.

8. Software and high performance computation. Reliable software for solving matrix equations has been available for a long time, due to its fundamental role in control applications; in particular, the SLICE Library was made available already in 1986. Early in the 1990s the SLICOT library (<http://www.slicot.org/>, [247]) replaced SLICE, and since then a large number of additions and improvements have been included; see, e.g., [230], [33]. Most recent versions of Matlab¹⁶ also rely on calls to SLICOT routines within the control-related Toolboxes. SLICOT includes a large variety of codes for model reduction and nonlinear problems on sequential and parallel architectures; as a workhorse, both the Bartels-Stewart algorithm and the Hessenberg-Schur algorithm are implemented. The Bartels-Stewart algorithm for triangular matrices is also included as a standard in LAPACK. Functions for solving the Lyapunov equation are also available in other computational environments, such as Mathematica¹⁷. Related projects have led to the developments of additional codes, which are usually available either in more general websites or directly from the authors. For instance, specifically oriented to linear matrix equations in the Matlab framework, the `lyapack` set of routines developed by Penzl in [198] has been particularly successful as a possible implementation of the ADI method for large scale Lyapunov equations, relying on preprocessing for the computation of quasi-optimal parameters. These routines were included in the NICONET Project repository (<http://www.icm.tu-bs.de/NICONET/>). The MESS set¹⁸ by Saak, Mena and Benner is the successor to the `lyapack` package, with the aim of fully exploiting the capabilities of newer releases of Matlab. In addition, MESS allows for the solution to a larger variety of matrix equations associated with the differential Riccati equation. A rather detailed list of routines for solving control-related matrix equations is provided in the book by Sima [221] and in the more recent book by Datta [64].

A number of benchmark problems have been made available for testing purposes. In addition to those available in the NICONET website, a variety of datasets is available in the Oberwolfach collection¹⁹, accompanied by a well documented description of the originating application problems; see also the description in [159].

Refined implementations of structured linear equation methods have been proposed for high performance computations. In particular, the efficient solution of triangular and quasi-triangular Sylvester equations has been discussed in [204], [200]. A high performance library for triangular Sylvester-type matrix equations (continuous and discrete-time) is also available at <http://www8.cs.umu.se/~isak/recsy/>, while a parallel SCALAPACK-style version of this software, called SCASY, is available at <http://www8.cs.umu.se/~granat/scasy.html>. Some of the SLICOT routines are overloaded in these libraries; see [141], [142], [98] for more information on the implementation on parallel architecture.

In [128] an early parallel algorithm was developed to solve medium size ($0 < n \leq 1000$) Lyapunov problems with a banded and negative definite matrix A ; experiments with a shared memory multiprocessor machine (Alliant FX-8) can also be found. The approach is similar in spirit to classical iterative linear system methods such as Jacobi

¹⁶MATLAB is a registered trademark of The MathWorks Inc.

¹⁷Mathematica is a registered trademark of Wolfram Research.

¹⁸Available at <http://www.en.mpi-magdeburg.mpg.de/mpcsc/mitarbeiter/saak/Software/mess.php?lang=en>

¹⁹Available at <http://portal.uni-freiburg.de/imteksimulation/downloads/benchmark>

and Gauss-Seidel. More recently, specialized parallel algorithms for Lyapunov, Stein and other generalized matrix equations for different modern architectures have been presented by a number of authors; see, e.g., [203] within the Cray T3E, [36],[37] employing a cluster of PCs, [23] within hybrid CPU-GPU platforms. The use of approaches based either on the square Smith iteration or on iterative techniques for the matrix sign function, as opposed to the Schur decomposition, is key to obtain good parallel performance.

Systems of matrix equations were implemented in a parallel environment in [45] and references therein. A parallel algorithm for the small scale solution to the multi-input Sylvester-observer equation (see section 7.2) was proposed in [43], and tested on two shared-memory vector machines.

9. Concluding remarks and future outlook. The solution of linear matrix equations has always attracted the attention of the engineering and scientific communities. The reliability of efficient core numerical linear algebra methods has made the solution of these matrix equations increasingly popular in application problem modelling. A good understanding of the theoretical tools and of the variety of numerical methods available for Sylvester-type equations provides a solid ground for attacking more general - nonlinear, multi-term or multi-functional - matrix equations. In particular, the efficient solution of multi-term matrix equations as those in (1.2) represents the next frontier for numerical linear algebra, as it currently is one of the major bottlenecks in the numerical treatment of PDEs involving stochastic terms. Advances in this direction will be tightly related to those which are being made in the solution of linear systems with tensor product structure, which in the simplest case can be written as

$$\mathbf{Ax} = b \quad \text{with} \quad \mathcal{A} = \sum_{j=1}^k I_{n_1} \otimes \cdots \otimes I_{n_{j-1}} \otimes A_j \otimes I_{n_{j+1}} \cdots \otimes I_{n_k}. \quad (9.1)$$

This problem is a further level of generalization of the standard Sylvester equation, where the solution is a k -way tensor, and its size explodes with k , even for modest values of n_j [153]. The complex Kronecker structure arising in (9.1) makes the problem very hard to even analyze, and its size calls for truncation or reduction procedures that rely on approximation theory and on hierarchical data structures; see, e.g., [104],[189],[190]. Scientific computing multi-dimensional applications can exploit these data tools to considerably lower the computational complexity of their model; see, e.g., [148] for a recent survey. Among the very recent projection methods used for the solution of (9.1), we find Krylov subspaces based procedures, that considerably generalize methods used for the two-dimensional case; see, e.g., [157],[9]. We envisage that a lot of scientific research will be devoted to multi-term and multi-dimensional problems in the forthcoming years.

We have limited our presentation mainly to linear problems. Non-linear matrix equations have a crucial and ever increasing role in many applications: for instance, the popular algebraic Riccati equation (see [164]) has a leading position in control applications and is an important tool in eigenvalue problems; we refer the reader to [42] for a very recent presentation of the rich literature on computational methods. Other fully nonlinear equations include, e.g, matrix eigenvalue problems ([80],[182]), and equations of the type $\mathbf{X} + A^\top F(\mathbf{X})A = Q$, where F is a properly defined nonlinear function of \mathbf{X} ; see, e.g., [205] and references therein.

Linear matrix equations with special properties arise when dealing with *periodic* dynamical systems. These problems give rise to periodic counterparts of the equa-

tions we have analyzed, such as Lyapunov and Sylvester equations. Corresponding Schur forms can be used for their solution, and necessary and sufficient conditions for a periodic discrete-time system to be equivalent to a time-invariant system are known; thorough treatments with developments on both the theoretical and algorithmic aspects, mainly on small size problems, have been carried out by R. Byers, P. Van Dooren, J. Sreedhar, A. Varga and their collaborators.

10. Acknowledgements. We are grateful to three reviewers for their very many helpful remarks. We would like to thank Ulrike Baur, Peter Benner, Michele Benzi, Lars Grasedyck, Yiding Lin, Daniel B. Szyld and Paul van Dooren, for helpful comments and pointers to many references. We are indebted with Daniel Szyld for providing us with [128]. We are grateful to Daniel Kressner and Froilán M. Dopico for carefully reading and amending a previous version of this manuscript, and for several helpful references. We are also grateful to the librarians of the Math Library at the Dipartimento di Matematica, Università di Bologna, for their professional dedication.

REFERENCES

- [1] M. Abramowitz and I. A. Stegun. *Handbook of Mathematical Functions*. Dover, 1965. [19, 37]
- [2] P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, Princeton, NJ, 2008. [42]
- [3] Mian Ilyas Ahmad, Imad Jaimoukha, and Michalis Frangos. Krylov subspace restart scheme for solving large-scale Sylvester equations. In *2010 American Control Conference*, Marriott Waterfront, Baltimore, MD, USA, June 30-July 02 2010. [24]
- [4] A. C. Antoulas. *Approximation of large-scale Dynamical Systems*. Advances in Design and Control. SIAM, Philadelphia, 2005. [5, 6, 19, 24, 42]
- [5] A. C. Antoulas, D. C. Sorensen, and Y. Zhou. On the decay rate of Hankel singular values and related issues. *Systems & Control Letters*, 46(5):323–342, August 2002. [29]
- [6] Ivo Babuska, Raul Tempone, and Georgios E. Zouraris. Galerkin finite element approximations of stochastic elliptic partial differential equations. *SIAM J. Numer. Anal.*, 42(2):800–825, 2004. [50]
- [7] Z. Bai and J. Demmel. Design of a parallel nonsymmetric eigenroutine toolbox, Part I. In R.F. Sincovec et al., editor, *Proc. of the 6th SIAM Conf. on Parallel Processing for Scientific Computing*, 1993. Also as Computer Science Division Tech Report UCB CSD-92-718, U. C. Berkeley, 1992. [31]
- [8] Zhaojun Bai and Daniel Skoogh. A projection method for model reduction of bilinear dynamical systems. *Linear Algebra and its Applications*, 415(2-3):406–425, 2006. Special Issue on Order Reduction of Large-Scale Systems. [50]
- [9] J. Ballani and L. Grasedyck. A projection method to solve linear systems in tensor format. *Num. Lin. Alg. Appl.*, 19:-, 2012. [53]
- [10] A. D. Barbour and S. Utev. Solving the Stein equation in compound Poisson approximation. *Adv. Appl. Probab.*, 30:449–475, 1998. [8, 42]
- [11] J. B. Barlow, M. M. Monahemi, and D. P. O’Leary. Constrained matrix Sylvester equations. *SIAM J. Matrix Anal. Appl.*, 13(1):1–9, 1992. [49]
- [12] S. Barnett. Simplification of the Lyapunov matrix equation $A^T P A - P = -Q$. *Automatic Control, IEEE Transactions on*, 19(4):446 – 447, Aug. 1974. [43]
- [13] A.Y. Barraud. A numerical algorithm to solve $A^T X A - X = Q$. *IEEE Trans. Automatic Control*, AC-22(5):883–885, October 1977. [43]

- [14] J. F. Barrett. Bibliography of A.M. Lyapunov's work. *Int. J. Control*, 55(3):785–790, 1992. [2]
- [15] R. H. Bartels and G. W. Stewart. Algorithm 432: Solution of the Matrix Equation $AX + XB = C$. *Comm. of the ACM*, 15(9):820–826, 1972. [3, 13, 14]
- [16] R. Bathia and P. Rosenthal. How and why to solve the operator equation $AX - XB = Y$. *Bull. London Math. Soc.*, 29:1–21, 1997. [2]
- [17] Ulrike Baur. *Control-Oriented Model Reduction for Parabolic Systems*. PhD thesis, Institute für Mathematik, TU Berlin, 2008. [40]
- [18] Ulrike Baur. Low rank solution of data-sparse Sylvester equations. *Numerical Linear Algebra with Applications*, 15(9):837–851, 2008. [27, 28]
- [19] Ulrike Baur and Peter Benner. Factorized solution of Lyapunov equations based on hierarchical matrix arithmetic. *Computing*, 78:211–234, 2006. [30, 40]
- [20] Ulrike Baur and Peter Benner. Cross-Gramian based model reduction for data-sparse systems. *Electronic Transactions on Numerical Analysis*, 31:256–270, 2008. [27]
- [21] Bernhard Beckermann. An Error Analysis for Rational Galerkin Projection Applied to the Sylvester Equation. *SIAM J. Numer. Anal.*, 49(6):2430–2450, 2011. [25, 26, 27, 33, 38]
- [22] P. Benner, T. Breiten, and T. Damm. Generalised tangential interpolation for model reduction of discrete-time MIMO bilinear systems. *International Journal of Control*, 84(8):1398–1407, 2011. [49]
- [23] P. Benner, P. Ezzatti, D. Kressner, E. S. Quintana-Ortí, and A. Rem'on. A mixed-precision algorithm for the solution of Lyapunov equations on hybrid CPU-GPU platforms. *Parallel Computing*, 37(8):439–450, 2011. [53]
- [24] P. Benner, P. Kürschner, and J. Saak. Efficient handling of complex shift parameters in the low-rank Cholesky factor ADI method. *Numerical Algorithms*, 62(2):225–251, 2013. [36]
- [25] P. Benner, R.-C. Li, and N. Truhar. On the ADI method for Sylvester equations. *Journal of Computational and Applied Mathematics*, 233(4):1035–1045, 2009. [26]
- [26] P. Benner, V. Mehrmann, and D. Sorensen (eds). *Dimension Reduction of Large-Scale Systems*. Lecture Notes in Computational Science and Engineering. Springer-Verlag, Berlin/Heidelberg, 2005. [5, 7, 41]
- [27] P. Benner and E. S. Quintana-Ortí. Solving stable generalized Lyapunov equations with the matrix sign function. *Numerical Algorithms*, 20(1):75–100, 1999. [30, 40, 46]
- [28] P. Benner and J. Saak. A Galerkin-Newton-ADI method for solving large-scale algebraic Riccati equations. Technical Report SPP1253-090, Deutsche Forschungsgemeinschaft - Priority Program 1253, 2010. [38, 39]
- [29] Peter Benner and Tobias Breiten. Low Rank Methods for a Class of Generalized Lyapunov Equations and Related Issues. *Numerische Mathematik*, 124(3):441–470, 2013. [50]
- [30] Peter Benner and Tobias Damm. Lyapunov equations, energy functionals, and model order reduction of bilinear and stochastic systems. *SIAM J. Control Optim.*, 49(2):686–711, 2011. [9, 49]
- [31] Peter Benner, Grece El Khoury, and Miloud Sadkane. On the squared Smith method for large-scale Stein equations. Technical Report MPIMD/12-15, Max Planck Institute Magdeburg, September 2012. [43, 44]
- [32] Peter Benner, Martin Köhler, and Jens Saak. Sparse-dense Sylvester equations in \mathcal{H}_2 -model order reduction. Technical Report MPIMD/11-11, Max Planck Institute Magdeburg Preprint, December 2011. Available from <http://www.mpi-magdeburg.mpg.de/preprints/>. [16]

- [33] Peter Benner, Daniel Kressner, Vasile Sima, and Andras Varga. The SLICOT Toolboxes - a survey. SLICOT Working Note 2009-1, August 2009. [14, 52]
- [34] Peter Benner, Patrick Krschner, and Jens Saak. A reformulated low-rank ADI iteration with explicit residual factors. *PAMM Proc. Appl. Math. Mech.*, 13:585–588, 2013. [36]
- [35] Peter Benner, Jing-Rebecca Li, and Thilo Penzl. Numerical solution of large-scale Lyapunov equations, Riccati equations, and linear-quadratic optimal control problems. *Num. Lin. Alg. with Appl.*, 15:1–23, 2008. [35]
- [36] Peter Benner, Enrique S. Quintana-Ortì, and Gregorio Quintana-Ortì. Solving stable Stein equations on distributed memory computers. In P. Amestoy, P. Berger, M. Daydè, I. Duff, V. Frayssè, L. Giraud, and D. Ruiz, editors, *EuroPar'99 Parallel Processing*, number 1685 in Lecture Notes in Computer Science, pages 1120–1123. Springer-Verlag, 1999. [53]
- [37] Peter Benner, Enrique S. Quintana-Ortì, and Gregorio Quintana-Ortì. Numerical solution of discrete stable linear matrix equations on multicomputers. *Parallel Algorithms and Applications*, 17(1):127–146, 2002. [53]
- [38] Peter Benner, Enrique S. Quintana-Ortì, and Gregorio Quintana-Ortì. Solving stable Sylvester equations via rational iterative schemes. *Journal of Scientific Computing*, 28(1):51–83, 2006. [15]
- [39] A. Betser, N. Cohen, and E. Zeheb. On solving the Lyapunov and Stein equations for a companion matrix. *Systems & Control Letters*, 25(3):211–218, June 1995. [29, 42]
- [40] J.H. Bevis, F.J. Hall, and R.E. Hartwig. The matrix equation $AX - XB = C$ and its special cases. *SIAM Journal on Matrix Analysis and Applications*, 9(3):348–359, 1988. [50]
- [41] W. G. Bickley and J. McNamee. Matrix and other direct methods for the solution of linear difference equation. *Philos. Trans. Roy. Soc. London Ser. A*, 252:69–131, 1960. [9, 15]
- [42] D.A. Bini, B. Iannazzo, and B. Meini. *Numerical Solution of Algebraic Riccati Equations*. SIAM, 2012. [7, 8, 53]
- [43] Christian H. Bischof, Biswa N. Datta, and Avijit Purkayastha. A Parallel Algorithm for the Multi-input Sylvester-Observer Equation. In *American Control Conference*, pages 152 – 156, 1992. [53]
- [44] Matthias Bollhöfer and André K. Eppler. A structure preserving FGMRES method for solving large Lyapunov equations. In M. Günther, A. Bartel, M. Brunk, S. Schoeps, and M. Striebel, editors, *Progress in Industrial Mathematics at ECMI 2010, Mathematics in Industry*, volume 17, pages 131–136, 2012. [40]
- [45] I. Borno. Parallel computation of the solutions of coupled algebraic Lyapunov equations. *Automatica*, 31(9):1345–1347, September 1995. [48, 53]
- [46] A. Bouhamidi and K. Jbilou. A note on the numerical approximate solutions for generalized Sylvester matrix equations with applications. *Applied Math. Comp.*, 206:687–694, 2008. [18, 45]
- [47] H. W. Braden. The equations $A^T X \pm X^T A = B$. *SIAM J. Matrix Anal. Appl.*, 20(2):295–302, 1998. [50]
- [48] Sanjoy Brahma and Biswa Datta. An optimization approach for minimum norm and robust partial quadratic eigenvalue assignment problems for vibrating structures. *Journal of Sound and Vibration*, 324(3-5):471–489, 2009. [7]
- [49] R. Byers. A LINPACK-style condition estimator for the equation $AX + XB^T = C$. *IEEE Trans. Automat. Control*, AC-29:926–928, 1984. [13]
- [50] R. Byers. Solving the algebraic Riccati equation with the matrix sign function. *Linear Algebra Appl.*, 85:267–279, 1987. [31]

- [51] R. Byers and D. Kressner. Structured condition numbers for invariant subspaces. *SIAM J. Matrix Anal. Appl.*, 28(2):326–347, 2006. [50]
- [52] D. Calvetti, B. Lewis, and L. Reichel. On the solution of large Sylvester-observer equations. *Numer. Linear Algebra Appl.*, 8:435–451, 2001. [48]
- [53] D. Calvetti and L. Reichel. Application of ADI iterative methods to the restoration of noisy images. *SIAM J. Matrix Anal. Appl.*, 17(1):165–186, January 1996. [7]
- [54] J. Carvalho and B. A. Datta. A block algorithm for the Sylvester-observer equation arising in state-estimation. In *Proceedings IEEE International Conference on Decision and Control*, pages 3898–3903, Orlando, 2001. [48]
- [55] J. Carvalho, K. Datta, and Yoopyo Hong. A new block algorithm for full-rank solution of the Sylvester-observer equation. *Automatic Control, IEEE Transactions on*, 48(12):2223–2228, Dec. 2003. [48]
- [56] Eugenio B. Castelan and Vilemar Gomes da Silva. On the solution of a Sylvester equation appearing in descriptor systems control theory. *Systems & Control Letters*, 54:109–117, February 2005. [49]
- [57] Chong Chen and Dan Schonfeld. Pose estimation from multiple cameras based on Sylvester’s equation. *Computer Vision and Image Understanding*, 114(6):652–666, June 2010. [8]
- [58] Marisa Condon and Rossen Ivanov. Nonlinear systems - algebraic Gramians and model reduction. *Compel*, 24(1):202–219, 2005. [49]
- [59] M. J. Corless and A. E. Frazho. *Linear systems and control - An operator perspective*. Pure and Applied Mathematics. Marcel Dekker, New York - Basel, 2003. [5]
- [60] O. L. V. Costa and C. S. Kubrusly. Lyapunov equation for infinite-dimensional discrete bilinear systems. *Systems & Control Letters*, 17(1):71–77, July 1991. [49]
- [61] T. Damm. Direct methods and ADI-preconditioned Krylov subspace methods for generalized Lyapunov equations. *Num. Lin. Alg. with Appl.*, 15:853–871, 2008. Special issue on Matrix equations. [49]
- [62] M. Darouach. Solution to Sylvester equation associated to linear descriptor systems. *Systems & Control Letters*, 55(10):835–838, October 2006. [49]
- [63] B. N. Datta and Y. Saad. Arnoldi methods for large Sylvester-like observer matrix equations, and an associated algorithm for partial spectrum assignment. *Lin. Alg. Appl.*, 154–156:225–244, 1991. [48, 49]
- [64] Biswa Nath Datta. *Numerical methods for linear control systems: design and analysis*. Elsevier, Academic Press, 2004. [52]
- [65] Biswa Nath Datta and Chandanie Hetti. Generalized Arnoldi methods for the Sylvester-observer equation and the multi-input pole placement problem. In *Proceedings of the 36th Conference on Decision & Control*, pages 4379–4383, San Diego, California USA, December 1997. [48]
- [66] B.N. Datta. Linear and numerical linear algebra in control theory: Some research problems. *Lin. Alg. Appl.*, 197–198:755–790, 1994. [5, 49]
- [67] E. de Souza and S. P. Bhattacharyya. Controllability, Observability and the Solution of $AX - XB = C$. *Lin. Alg. Appl.*, 39:167–188, 1981. [12]
- [68] J. Demmel. Three methods for refining estimates of invariant subspaces. *Computing*, 38:43–57, 1987. [7]
- [69] Feng Ding and Tongwen Chen. Iterative least-squares solutions of coupled Sylvester matrix equations. *Systems & Control Letters*, 54(2):95–107, February 2005. [48]
- [70] P. Van Dooren. Reduced order observers: a new algorithm and proof. *System and Control Letters*, 4:243–251, 1984. [48]

- [71] F. M. Dopico, J. Gonzalez, D. Kressner, and V. Simoncini. Projection methods for large T -Sylvester equations. Technical report, Università di Bologna, 2014. in preparation. [50, 51]
- [72] V. Druskin and L. Knizhnerman. Extended Krylov subspaces: approximation of the matrix square root and related functions. *SIAM J. Matrix Anal. Appl.*, 19(3):755–771, 1998. [23]
- [73] V. Druskin and V. Simoncini. Adaptive rational Krylov subspaces for large-scale dynamical systems. *Systems and Control Letters*, 60:546–560, 2011. [33]
- [74] V. Druskin and V. Simoncini. Adaptive rational Krylov subspaces for large-scale dynamical systems. *Systems and Control Letters*, 60:546–560, 2011. [33, 34, 38]
- [75] V. Druskin, V. Simoncini, and M. Zaslavsky. Adaptive tangential interpolation in rational Krylov subspaces for MIMO model reduction data. Technical report, Dipartimento di Matematica, Università di Bologna, November 2012. [34, 39]
- [76] Vladimir Druskin, Leonid Knizhnerman, and V. Simoncini. Analysis of the rational Krylov subspace and ADI methods for solving the Lyapunov equation. *SIAM J. Numer. Anal.*, 49:1875–1898, 2011. [33, 37, 38]
- [77] Guangren Duan. A note on combined generalized Sylvester matrix equations. *Journal of Control Theory and Applications*, 4:397–400, 2004. [47, 48]
- [78] Nancy S. Ellner and Eugene L. Wachspress. New ADI model problem applications. In *Proceedings of 1986 ACM Fall joint computer conference, Dallas, Texas, United States*, pages 528–534. IEEE Computer Society Press, Los Alamitos, CA, 1986. [8, 19, 25]
- [79] Nancy S. Ellner and Eugene L. Wachspress. Alternating Direction Implicit iteration for systems with complex spectra. *SIAM J. Numer. Anal.*, 23(3):859–870, 1991. [19, 37]
- [80] Howard C. Elman and Minghao Wu. Lyapunov inverse iteration for computing a few rightmost eigenvalues of large generalized eigenvalue problems. *SIAM. J. Matrix Anal. Appl.*, 34(4):16851707, 2013. [6, 7, 53]
- [81] E. Emre and L. M. Silverman. The equation $xr + qy = \phi$: A characterization of solutions. *SIAM J. Control Optim.*, 19(1):33–38, 1981. [46]
- [82] M. Epton. Methods for the solution of $ADXD - BXC = E$ and its application in the numerical solution of implicit ordinary differential equations. *BIT*, 20:341–345, 1980. [8]
- [83] O.G. Ernst, C.E. Powell, D.J. Silvester, and E. Ullmann. Efficient solvers for a linear stochastic Galerkin mixed formulation of diffusion problems with random data. *SIAM J. Sci. Comput.*, 31:1424–1447, 2009. [9, 10, 47]
- [84] K. V. Fernando and H. Nicholson. On the structure of balanced and other principal representations of SISO systems. *Automatic Control, IEEE Transactions on*, 28(2):228–231, 1983. [6]
- [85] K. V. Fernando and H. Nicholson. On a fundamental property of the cross-Gramian matrix. *IEEE Trans. Circ. Syst.*, CAS-31(5):504–505, 1984. [6]
- [86] G. Flagg and S. Gugercin. On the ADI method for the Sylvester equation and the optimal-H2 points. *Applied Numerical Mathematics*, 64:50–58, 2013. [26]
- [87] Darran Furnival, Howard Elman, and Catherine Powell. H(div) preconditioning for a mixed finite element formulation of the stochastic diffusion problem. *Mathematics of Computation*, 79:733–760, 2010. [10]
- [88] Z. Gajić and M. Qureshi. *Lyapunov Matrix Equation in System Stability and Control*. Math. in Science and Engineering. Academic Press, San Diego, CA, 1995. [2, 5, 7, 12, 28, 29, 30]

- [89] K. Gallivan, A. Vandendorpe, and P. Van Dooren. Sylvester equations and projection-based model reduction. *J. Comput. Appl. Math.*, 162(1):213–229, 2004. [6]
- [90] Stephan Ramon Garcia and Amy L. Shoemaker. On the matrix equation $AX + XA^T = 0$. *Lin. Alg. Appl.*, 438:2740–2746, 2013. In Press. [52]
- [91] Judith D. Gardiner, Alan J. Laub, James J. Amato, and Cleve B. Moler. Solution of the Sylvester matrix equation $AXB^T + CXD^T = E$. *ACM Transactions on Mathematical Software*, 18(2):223–231, June 1992. [45]
- [92] Judith D. Gardiner, Matthew R. Wette, Alan J. Laub, James J. Amato, and Cleve B. Moler. Algorithm 705: A FORTRAN-77 software package for solving the Sylvester matrix equation $AXB^T + CXD^T = E$. *ACM Transactions on Mathematical Software*, 18(2):232–238, June 1992. [45]
- [93] Ali R. Ghavimi and Alan J. Laub. Numerical methods for nearly singular constrained matrix Sylvester equations. *SIAM J. Matrix Anal. Appl.*, 17:212–221, 1996. [49]
- [94] I. Gohberg, T. Kailath, and V. Olshevsky. Fast Gaussian elimination with partial pivoting for matrices with displacement structure. *Math. Comp.*, 64(212):1557–1576, 1995. [15]
- [95] I. Gohberg, P. Lancaster, and L. Rodman. *Invariant Subspaces of Matrices with Applications*. J. Wiley & Sons, New York, 1986. [11]
- [96] G. Golub and C. F. Van Loan. *Matrix Computations*. The Johns Hopkins University Press, Baltimore, 3rd edition, 1996. [12, 14]
- [97] G. H. Golub, S. Nash, and C. Van Loan. A Hessenberg-Schur method for the problem $AX + XB = C$. *IEEE Trans. Autom. Contr.*, AC-24(6):909–913, 1979. [14]
- [98] Robert Granat. Parallel algorithms and condition estimators for standard and generalized triangular Sylvester-type matrix equations. In *PARA06 - State of the Art in Scientific and Parallel Computing. Vol. 4699. Lecture Notes in Computer Science*. Springer, 2007. [52]
- [99] L. Grasedyck. Existence of a low rank or \mathcal{H} -matrix approximant to the solution of a Sylvester equation. *Numerical Linear Algebra with Appl.*, 11(4):371–389, 2004. [28]
- [100] L. Grasedyck. Nonlinear multigrid for the solution of large-scale Riccati equations in low-rank and \mathcal{H} -matrix format. *Numer. Linear Algebra Appl.*, 15:779–807, 2008. [40]
- [101] L. Grasedyck and W. Hackbusch. A multigrid method to solve large scale Sylvester equations. *SIAM J. Matrix Anal. Appl.*, 29:870–894, 2007. [27, 28]
- [102] L. Grasedyck, W. Hackbusch, and B. Khoromskij. Solution of large scale algebraic matrix Riccati equations by use of hierarchical matrices. *Computing*, 70(2):121–165, 2003. [40]
- [103] Lars Grasedyck. *Theorie und Anwendungen Hierarchischer Matrizen*. PhD thesis, Mathematisch-Naturwissenschaftliche Fakultät, Christian-Albrechts-Universität zu Kiel, 2001. [27]
- [104] Lars Grasedyck, Daniel Kressner, and Christine Tobler. A literature survey of low-rank tensor approximation techniques. Technical Report 1302.7121v1, Arxiv, February 2013. [53]
- [105] W. S. Gray and J. Mesko. Energy functions and algebraic Gramians for bilinear systems. In *Preprints of the 4th IFAC Nonlinear Control Systems Design Symposium*, Enschede, The Netherlands, 1998. [9, 49]
- [106] E. Grimme. *Krylov projection methods for model reduction*. PhD thesis, The University of Illinois at Urbana-Champaign, 1997. [24]
- [107] T. Gudmundsson and A. J. Laub. Approximate solution of large sparse Lyapunov equations. *IEEE Trans. Automatic Control*, 39(5):1110–1114, 1994. [41]

- [108] A. El Guennouni, K. Jbilou, and A.J. Riquet. Block Krylov subspace methods for solving large Sylvester equations. *Numerical Algorithms*, 29:75–96, 2002. [16, 18]
- [109] S. Gugercin, A. C. Antoulas, and C. Beattie. \mathcal{H}_2 model reduction for large-scale linear dynamical systems. *SIAM J. Matrix Anal. Appl.*, 30:609–638, 2008. [38]
- [110] S. Gugercin, D. C. Sorensen, and A. C. Antoulas. A modified low-rank Smith method for large-scale Lyapunov equations. *Numerical Algorithms*, 32:27–55, 2003. [36]
- [111] M. H. Gutknecht. Krylov subspace algorithms for systems with multiple right hand sides: a collection of references, 2004. available at <http://www.sam.math.ethz.ch/~mhg/talks/abstr/abstr.html>. [4, 22]
- [112] Stefan Güttel. Rational Krylov approximation of matrix functions: Numerical methods and optimal pole selection. *GAMM Mitteilungen*, 36(1):8–31, 2013. [24, 25]
- [113] S. J. Hammarling. Numerical solution of the stable, non-negative definite Lyapunov equation. *IMA J. Numer. Anal.*, 2:303–323, 1982. [30, 46]
- [114] Sven Hammarling. Numerical solution of the discrete-time, convergent, non-negative definite Lyapunov equation. *Systems & Control Letters*, 17(2):137–139, August 1991. [43]
- [115] B. Hanzon and R. L. M. Peeters. A Faddeev sequence method for solving Lyapunov and Sylvester equations. *Lin. Alg. Appl.*, 241–243:401–430, 1996. [51]
- [116] C. Hartmann, B. Schäfer-Bung, and A. Zueva. Balanced averaging of bilinear systems with applications to stochastic control. *SIAM J. Control Optim.*, 51:2356–2378, 2013. [9]
- [117] He Kong and Bin Zhou and Mao-Rui Zhang. A Stein equation approach for solutions to the Diophantine equations. In *Control and Decision Conference (CCDC), 2010 Chinese*, pages 3024 – 3028, 2010. [42, 51]
- [118] J. Z. Hearon. Nonsingular solutions of $TA - BT = C$. *Lin. Alg. Appl.*, 16:57–63, 1977. [12]
- [119] D. Heinig and U. Jungnickel. Lyapunov equations for companion matrices. *Lin. Alg. Appl.*, 76:137–147, 1986. [29]
- [120] G. Hoyer and C. Kenney. The sensitivity of the stable Lyapunov equation. *SIAM J. Control and Optimization*, 26:321–344, 1988. [28]
- [121] M. Heyouni and K. Jbilou. Matrix Krylov subspace methods for large scale model reduction problems. *Applied Mathematics and Computation*, 181:1215–1228, October 2006. [35]
- [122] N. J. Higham. Perturbation theory and backward error for $AX - XB = C$. *BIT*, 33:124–136, 1993. [13, 28]
- [123] Nicholas J. Higham. *Accuracy and Stability of Numerical Algorithms*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1996. [13]
- [124] Nicholas J. Higham. *Matrix Functions – Theory and Applications*. SIAM, Philadelphia, USA, 2008. [25]
- [125] M. Hochbruck and G. Starke. Preconditioned Krylov subspace methods for Lyapunov matrix equations. *SIAM Matrix Anal. and Appl.*, 16(1):156–171, 1995. [39]
- [126] A. Hodel, K. Poolla, and B. Tenison. Numerical solution of the Lyapunov equation by approximate power iteration. *Lin. Alg. Appl.*, 236:205–230, 1996. [41]
- [127] A. Scottedward Hodel and Pradeep Misra. Least-squares approximate solution of overdetermined Sylvester equations. *SIAM J. Matrix Anal. Appl.*, 18(2):279–290, April 1997. [21]

- [128] Alan Scottedward Hodel. *Numerical methods for the solution of large and very large, sparse Lyapunov equations*. PhD thesis, Engineering, Electronics and Electrical, University of Illinois at Urbana-Champaign, 1989. [31, 52, 54]
- [129] A.Scottedward Hodel and K. Poolla. Numerical solution of very large, sparse Lyapunov equations through approximate power iteration. In *Proceedings of the 29th Conference on Decision and Control, IEEE, WA12*, pages 291–296, December 1990. [41]
- [130] R. A. Horn and C. R. Johnson. *Topics in Matrix Analysis*. Cambridge University Press, Cambridge, 1991. [2, 5, 11, 28, 47]
- [131] D. Y. Hu and L. Reichel. Krylov Subspace Methods for the Sylvester Equation. *Linear Algebra and Appl.*, 172:283–313, July 1992. [21, 23, 24, 31, 35]
- [132] Guang-Da Hu and Qiao Zhu. Bounds of modulus of eigenvalues based on Stein equation. *Kybernetika*, 46(4):655–664, 2010. [42]
- [133] Kh. D. Ikramov. Conditions for unique solvability of the matrix equation $AX + X^T B = C$. *Doklady Mathematics*, 81(1):63–65, 2010. [50]
- [134] I. M. Jaimoukha and E. M. Kasenally. Oblique projection methods for large scale model reduction. *SIAM J. Matrix Anal. Appl.*, 16(2):602–627, 1995. [23, 35]
- [135] I. M. Jaimoukha and E. M. Kasenally. Krylov subspace methods for solving large Lyapunov equations. *SIAM J. Numer. Anal.*, 31(1):227–251, Feb. 1994. [32, 35]
- [136] Imad M. Jaimoukha and Ebrahim M. Kasenally. Implicitly Restarted Krylov Subspace Methods for Stable Partial Realizations. *SIAM J. Matrix Anal. Appl.*, 18(3):633–652, 1997. [35]
- [137] K. Jbilou. Low rank approximate solutions to large Sylvester matrix equations. *Applied Mathematics and Computation*, 177:365–376, 2006. [23, 43]
- [138] K. Jbilou. ADI preconditioned Krylov methods for large Lyapunov matrix equations. *Lin. Alg. Appl.*, 432:2473–2485, May 2010. [39]
- [139] K. Jbilou, A. Messaoudi, and H. Sadok. Global FOM and GMRES algorithms for matrix equations. *Applied Numerical Mathematics*, 31:49–63, 1999. [23]
- [140] K. Jbilou and A.J. Riquet. Projection methods for large Lyapunov matrix equations. *Linear Algebra and its Applications*, 415(2-3):344–358, 2006. [34]
- [141] I. Jonsson and B. Kågström. Recursive blocked algorithms for solving triangular systems Part I: One-sided and coupled Sylvester-type matrix equations. *ACM Trans. Math. Softw.*, 28(4):392–415, 2002. [14, 15, 52]
- [142] I. Jonsson and B. Kågström. Recursive blocked algorithms for solving triangular systems Part II: Two-sided and generalized Sylvester and Lyapunov matrix equations. *ACM Trans. Math. Softw.*, 28(4):416–435, 2002. [14, 52]
- [143] K. Jbilou and A. Messaoudi. A Computational Method for Symmetric Stein Matrix Equations. In *Numerical Linear Algebra in Signals, Systems and Control*, volume 80 of *Lecture Notes in Electrical Engineering*, pages 295–311. Springer, 2011. [43]
- [144] B. Kågström and P. Van Dooren. A generalized state-space approach for the additive decomposition of a transfer matrix. *Int. J. Numerical Linear Algebra with Applications*, 1(2):165–181, 1992. [47]
- [145] B. Kågström and P. Poromaa. LAPACK-Style algorithms and software for solving the generalized Sylvester equation and estimating the separation between regular matrix pairs. *ACM Transactions on Mathematical Software*, 22(1):78–103, 1996. Also as LAPACK Working Note 75. [48]
- [146] Bo Kågström and Lars Westin. Generalized Schur methods with Condition Estimators for Solving the Generalized Sylvester equation. *Automatic Control, IEEE Transactions on*, 34:745–751, 1989. [47, 48]

- [147] Thomas Kailath. *Linear systems*. Englewood Cliffs, N.J.; Prentice-Hall, 1980. [43]
- [148] B. N. Khoromskij. Tensors-structured numerical methods in scientific computing: survey on recent advances. *Chemometrics and Intelligent Laboratory systems*, 110:1–19, 2012. [53]
- [149] A. Klein and P. Spreij. On Fisher’s information matrix of an ARMAX process and Sylvester’s resultant matrices. *Linear Algebra Appl.*, 237/238:579–590, 1996. [8]
- [150] A. Klein and P.J.C. Spreij. On Stein’s equation, Vandermonde matrices and Fisher’s information matrix of time series processes. Part I: The autoregressive moving average process. *Lin. Alg. Appl.*, 329(1-3):9–47, 2001. [8, 42]
- [151] André Klein and Peter Spreij. On the solution of Stein’s equation and Fisher’s information matrix of an ARMAX process. *Linear Algebra and its Applications*, 396:1–34, 2005. [8, 42]
- [152] L. Knizhnerman and V. Simoncini. Convergence analysis of the Extended Krylov Subspace Method for the Lyapunov equation. *Numerische Mathematik*, 118(3):567–586, 2011. [33]
- [153] T. G. Kolda and B. W. Bader. Tensor decompositions and applications. *SIAM Rev.*, 51:455–500, 2009. [53]
- [154] N. Komaroff. Simultaneous eigenvalue lower bounds for the Lyapunov matrix equation. *IEEE Trans. on Autom. Control*, 33(1):126–128, 1988. [29]
- [155] M. Konstantinov, D. Gu, V. Mehrmann, and P. Petkov. *Perturbation theory for matrix equations*. Studies in Computational Mathematics 9. Elsevier, 2003. [3, 12, 13, 47]
- [156] D. Kressner. Block variants of Hammarling’s method for solving Lyapunov equations. *ACM Trans. Math. Software*, 34(1):1–15, 2008. [30, 43]
- [157] D. Kressner and C. Tobler. Krylov subspace methods for linear systems with tensor product structure. *SIAM J. Matrix Anal. Appl.*, 31(4):16881714, 2010. [53]
- [158] Daniel Kressner. Memory-efficient Krylov subspace techniques for solving large-scale Lyapunov equations. In *IEEE International Symposium on Computer-Aided Control Systems*, pages 613–618, San Antonio, 2008. [25]
- [159] Daniel Kressner, Volker Mehrmann, and Thilo Penzl. CTLEX - a collection of benchmark examples for continuous-time Lyapunov equations. SLICOT Working Note 1999-6, May 1999. [52]
- [160] Daniel Kressner, Christian Schröder, and David S. Watkins. Implicit QR algorithms for palindromic and even eigenvalue problems. *Numer. Algor.*, 51:209–238, 2009. [50]
- [161] Daniel Kressner and Christine Tobler. Low-rank tensor Krylov subspace methods for parametrized linear systems. *SIAM. J. Matrix Anal. & Appl.*, 32(4):1288–1316, 2011. [50]
- [162] Peter Kunkel and Volker Ludwig Mehrmann. *Differential-Algebraic Equations: Analysis and Numerical Solution*. European Mathematical Society, 2006. [3, 44]
- [163] P. Lancaster, L. Lerer, and M. Tismenetsky. Factored forms for solutions of $AX - XB = C$ and $X - AXB = C$ in companion matrices. *Lin. Alg. Appl.*, 62:19–49, 1984. [42]
- [164] P. Lancaster and L. Rodman. *Algebraic Riccati equations*. Oxford Univ. Press, 1995. [42, 53]
- [165] P. Lancaster and M. Tismenetsky. *The Theory of Matrices*. Academic Press, New York, second edition, 1985. [11, 28, 42]
- [166] Peter Lancaster. Explicit solutions of linear matrix equations. *SIAM Review*, 12(4):544–566, 1970. [2, 3, 11, 12]

- [167] V. Larin and F. Aliev. Construction of square root factor for solution of a Lyapunov matrix equation. *Systems Control Lett.*, 20:109–112, 1993. [30]
- [168] A. Laub. Numerical linear algebra aspects of control design computations. *IEEE Trans. Automat. Control*, AC-30:97–108, 1985. [7]
- [169] R. B. Lehoucq, D. C. Sorensen, and C. Yang. *ARPACK Users Guide: Solution of large scale Eigenvalue Problems by Implicitly Restarted Arnoldi Methods*. SIAM, Philadelphia, 1998. [41]
- [170] L. Lerer and A.C.M. Ran. A new inertia theorem for Stein equations, inertia of invertible Hermitian block Toeplitz matrices and matrix orthogonal polynomials. *Integral Equations and Operator Theory*, 47:339–360, 2003. [42]
- [171] J.-R. Li and J. White. Low-Rank solutions of Lyapunov equations. *SIAM J. Matrix Anal. Appl.*, 24(1):260–280, 2002. [36, 38]
- [172] Jing-Rebecca Li. *Model Reduction of Large Linear Systems via Low Rank System Gramians*. PhD thesis, MIT, 2000. [31]
- [173] Tiexiang Li, Peter Chang-Yi Weng, Eric King wah Chu, and Wen-Wei Lin. Large-scale Stein and Lyapunov equations, Smith method, and applications. *Numer. Algor.*, 63:727–752, 2013. [43]
- [174] Joerg Liesen and Zdenek Strakos. *Krylov Subspace Methods. Principles and analysis*. Oxford University Press, 2012. [29]
- [175] Matthew M. Lin and Chun-Yueh Chiang. On the T-Stein matrix equation $X = AX^T B + C$. Technical Report 2012-10-002, National Center for Theoretical Sciences, National Tsing Hua University, Taiwan, 2012. [44]
- [176] Y. Lin and V. Simoncini. Minimal residual methods for large scale Lyapunov equations. *Applied Num. Math.*, 72:52–71, 2013. [21, 31, 35]
- [177] Yiqin LIN, Liang BAO, Yanhua CAO, and Liping ZHOU. A direct method for solving projected generalized continuous-time sylvester equations. *J. Control Theory Appl.*, 10(3):319–324, 2012. [45]
- [178] Yiqin Lin and Yimin Wei. Condition numbers of the generalized Sylvester equation. *IEEE Transactions on Automatic Control*, 52:2380 – 2385, Dec. 2007. [47]
- [179] An Lu and Eugene L. Wachspress. Solution of Lyapunov equations by Alternating Direction Implicit iteration. *Computers Math. Applic.*, 21(9):43–58, 1991. [35]
- [180] D. Luenberger. Observers for multivariable systems. *IEEE Trans. Automat. Control*, 11:190–197, 1966. [7, 48]
- [181] M. Mariton. *Jump linear systems in automatic control*. Marcel Dekker, New York - Basel, 1990. [9]
- [182] Karl Meerbergen and Alastair Spence. Inverse iteration for purely imaginary eigenvalues with application to the detection of Hopf bifurcations in large scale problems. *SIAM Journal on Matrix Analysis and Applications*, 31(4):1982–1999, May 2010. [6, 7, 53]
- [183] Carl Christian Kjelgaard Mikkelsen. *Numerical methods for large Lyapunov equations*. PhD thesis, Purdue University, West Lafayette, Indiana, May 2009. [31, 39, 40]
- [184] C. B. Moler and G. W. Stewart. An algorithm for generalized eigenvalue problems. *SIAM J. Numer. Anal.*, 10:241–256, 1973. [46]
- [185] Moghen M. Monahemi, Jewel B. Barlow, and Dianne P. O’Leary. The design of reduced order observers with precise loop transfer recovery. *AIAA Journal of Guidance, Control, and Dynamics*, 15(6):1320–1326, 1992. [49]
- [186] C.T. Mullis and R.A. Roberts. Synthesis of minimum roundoff noise fixed-point digital filters. *IEEE Trans. on Circuits and Systems*, CAS-23(9):551–562, September 1976. [6]

- [187] Hung Dinh Nong. *Numerical Solutions of Matrix Equations Arising in Model Reduction of Large-Scale Linear-Time-Invariant Systems*. PhD thesis, Electrical and Computer Engineering, Rice University, January 2010. [31, 41]
- [188] Ryan Nong and Danny C. Sorensen. A parameter free ADI-like method for the numerical solution of large scale Lyapunov equations. Technical report, CAM - Rice University, May 2009. [41]
- [189] I. V. Oseledets and E. E. Tyrtyshnikov. Breaking the curse of dimensionality, or how to use SVD in many dimensions. *SIAM Journal on Scientific Computing*, 31(5):3744–3759, 2009. [53]
- [190] I. V. Oseledets and E. E. Tyrtyshnikov. TT-Cross approximation for multidimensional arrays. Technical Report 05, Institute of Numerical Mathematics, Russian Academy of Sciences, June 2009. [53]
- [191] R. Patel, A. J. Laub, and P. Van Dooren. *Numerical Linear Algebra Techniques for Systems and Control*. IEEE Press, Piscataway NJ, 1993. [7]
- [192] D. W. Peaceman and Jr. H. H. Rachford. The numerical solution of parabolic and elliptic equations. *Journal of SIAM*, 3:28–41, 1955. [25]
- [193] Ralf Peeters and Paolo Rapisarda. A two-variable approach to solve the polynomial Lyapunov equation. *Systems & Control Letters*, 42(2):117–126, February 2001. [51]
- [194] T. Penzl. Numerical solution of generalized Lyapunov equations. *Advances in Computational Mathematics*, 8:33–48, 1998. [45, 46]
- [195] T. Penzl. *Numerische Lösung großer Lyapunov-Gleichungen*. Logos-Verlag, Berlin, 1998. Diss. Fakultät für Mathematik, TU Chemnitz. [31]
- [196] T. Penzl. A cyclic low-rank Smith method for large sparse Lyapunov equations. *SIAM J. Sci. Comput.*, 21(4):1401–1418, 2000. [35, 36, 37, 38, 43, 44]
- [197] T. Penzl. Eigenvalue decay bounds for solutions of Lyapunov equations: the symmetric case. *Systems and Control Letters*, 40(2):139–144, 2000. [19, 29]
- [198] T. Penzl. Lyapack users guide. Technical Report SFB393/00-33, TU Chemnitz, 09107 Chemnitz, D, 2000. Available from <http://www.tu-chemnitz.de/sfb393/sfb00pr.html>. [35, 38, 52]
- [199] V. Popov. Hyperstability and optimality of automatic systems with several control functions. *Rev. Roum. Sci. Tech.*, 9:629–690, 1964. [43]
- [200] Peter Poromaa. Parallel algorithms for triangular Sylvester equations: Design, scheduling and scalability issues. In *Applied Parallel Computing Large Scale Scientific and Industrial Problems*, volume 1541, pages 438–446. Springer, 1998. [52]
- [201] C. Powell and D. Silvester. Preconditioning steady-state Navier-Stokes equations with random data. *SIAM Journal Sci. Comp.*, 34(5):A2482–A2506, 2012. [10]
- [202] H.M. Power. A note on the matrix equation $A'LA - L = -K$. *Automatic Control, IEEE Transactions on*, 14(4):411–412, August 1969. [42]
- [203] E. S. Quintana-Ortí and R. van de Geijn. Specialized parallel algorithms for solving Lyapunov and Stein equations. *J. of Parallel and Distributed Computing*, 61(10):1489–1504, 2001. [53]
- [204] E. S. Quintana-Ortí and R. van de Geijn. Formal derivation of algorithms: the triangular Sylvester equation. *ACM Trans. on Mathematical Software*, 29:218–243, 2003. [52]
- [205] A. C. M. Ran, M. C. B. Reurings, and L. Rodman. A perturbation analysis for nonlinear selfadjoint operator equations. *SIAM J. Matrix Anal. Appl.*, 28:89–104, 2006. [53]

- [206] Hongwu Ren and Richard Dekany. Fast wave-front reconstruction by solving the Sylvester equation with the alternating direction implicit method. *Optics Express*, 12:3279–3296, 2004. [8]
- [207] Mickaël Robbé and Miloud Sadkane. A convergence analysis of GMRES and FOM methods for Sylvester equations. *Numerical Algorithms*, 30:71–89, 2002. [17, 18]
- [208] J. D. Roberts. Linear model reduction and solution of the algebraic Riccati equation by use of the sign function. *Int. J. Control*, 32(32):677–687, 1980. [30]
- [209] William E. Roth. The equations $AX - YB = C$ and $AX - XB = C$ in matrices. *Proceedings of the American Mathematical Society*, 3(3):392–396, Jun. 1952. [10, 46, 47]
- [210] A. Ruhe. Rational Krylov sequence methods for eigenvalue computation. *Lin. Alg. Appl.*, 58:391–405, 1984. [24]
- [211] Y. Saad. Numerical solution of large Lyapunov equations. In M. A. Kaashoek, J. H. van Schuppen, and A. C. Ran, editors, *Signal Processing, Scattering, Operator Theory, and Numerical Methods. Proceedings of the international symposium MTNS-89, vol III*, pages 503–511, Boston, 1990. Birkhauser. [22, 32]
- [212] Y. Saad. *Iterative methods for sparse linear systems*. SIAM, Society for Industrial and Applied Mathematics, 2nd edition, 2003. [4, 16, 18, 24]
- [213] John Sabino. *Solution of Large-Scale Lyapunov Equations via the Block Modified Smith Method*. PhD thesis, Rice University, 2006. [19, 29, 31]
- [214] Miloud Sadkane. A low-rank Krylov squared Smith method for large-scale discrete-time Lyapunov equations. *Linear Algebra and its Applications*, 436(8):2807–2827, 2012. [43, 44]
- [215] B. Schäfer-Bung, C. Hartmann, B. Schmidt, and C. Schütte. Dimension reduction by balanced truncation: Application to light-induced control of open quantum systems. *The Journal of Chemical Physics*, 135:014112, 2011. (13 pages). [9]
- [216] W. H. A. Schilders, H. A. van der Vorst, and J. Rommes. *Model Order Reduction: Theory, Research Aspects and Applications*. Springer-Verlag, Berlin/Heidelberg, 2008. [5, 7]
- [217] Peter J. Schmid. Nonmodal stability theory. *Annu. Rev. Fluid Mech.*, 39:129–162, 2007. [7]
- [218] S. Shank and V. Simoncini. Krylov subspace methods for large scale constrained Sylvester equations. *SIAM J. Matrix Anal. Appl.*, 34(4):1448–1463, 2013. [49]
- [219] Fernando C. Silva and Rita Simões. On the Lyapunov and Stein Equations, I. *Linear Algebra and its Applications*, 420(23):329–338, 2007. [28, 42]
- [220] Fernando C. Silva and Rita Simões. On the Lyapunov and Stein Equations, II. *Linear Algebra and its Applications*, 426(23):305–311, 2007. [28, 42]
- [221] Vasile Sima. *Algorithms for linear-quadratic optimization*. Marcel Dekker, Inc., 1996. [14, 30, 52]
- [222] Vasile Sima and Peter Benner. Experimental evaluation of new SLICOT solvers for linear matrix equations based on the matrix sign function. *IEEE Multi-conference on Systems and Control*, pages 601–606, 2008. 9th IEEE International Symposium on Computer-Aided Control Systems Design (CACSD). [40]
- [223] V. Simoncini. On the numerical solution of $AX - XB = C$. *BIT*, 36(4):814–830, 1996. [12, 17, 21, 23]
- [224] V. Simoncini. A new iterative method for solving large-scale Lyapunov matrix equations. *SIAM J. Sci. Comput.*, 29(3):1268–1288, 2007. [32, 38]

- [225] V. Simoncini. The Extended Krylov subspace for parameter dependent systems. *Applied Numer. Math.*, 60(5):550–560, 2010. [19]
- [226] V. Simoncini and V. Druskin. Convergence analysis of projection methods for the numerical solution of large Lyapunov equations. *SIAM J. Numer. Anal.*, 47(2):828–843, 2009. [32]
- [227] V. Simoncini and M. Sadkane. Arnoldi-Riccati method for large eigenvalue problems. *BIT Numerical Mathematics*, 36(3):579–594, 1996. [7]
- [228] V. Simoncini and D. B. Szyld. Recent computational developments in Krylov subspace methods for linear systems. *Num. Lin. Alg. with Appl.*, 14(1):1–59, 2007. [16, 39]
- [229] John R. Singler. Approximate low rank solutions of Lyapunov equations via Proper Orthogonal Decomposition. In *2008 American Control Conference, Westin Seattle Hotel, Seattle, Washington, USA*, June 2008. WeA08.3. [41]
- [230] Martin Slowik, Peter Benner, and Vasile Sima. Evaluation of the linear matrix equation solvers in SLICOT. *J. of Numerical Analysis, Industrial and Applied Mathematics*, 2(1-2):11–34, 2007. [14, 52]
- [231] R. A. Smith. Matrix equation $XA + BX = C$. *SIAM J. Appl. Math.*, 16(1):198–201, 1968. [25, 43]
- [232] D. Sorensen and Y. Zhou. Bounds on eigenvalue decay rates and sensitivity of solutions to Lyapunov equations. Technical Report 02-07, Rice University, Houston, Texas, 2002. [29]
- [233] D. C. Sorensen and Y. Zhou. Direct methods for matrix Sylvester and Lyapunov equations. *J. Appl. Math.*, 2003:277–303, 2003. [14]
- [234] D.C. Sorensen and A.C. Antoulas. The Sylvester equation and approximate balanced reduction. *Lin. Alg. Appl.*, 351–352:671–700, 2002. [6, 16]
- [235] G. Starke and W. Niethammer. SOR for $AX - XB = C$. *Lin. Alg. Appl.*, 154-156:355–375, 1991. [27]
- [236] G. W. Stewart and J-G. Sun. *Matrix Perturbation Theory*. Academic Press, 1990. [7, 12, 13, 45, 46, 47, 48]
- [237] T. Stykel. Numerical solution and perturbation theory for generalized lyapunov equations. *Linear Algebra Appl.*, 349:155–185, 2002. [45]
- [238] J. J. Sylvester. Sur l'equations en matrices $px = xq$. *Comptes Rendus Acad. Sci. Paris*, 99:67–71,115–116, 1884. [2]
- [239] V.L. Syrmos. Disturbance decoupling using constrained Sylvester equations. *IEEE Transactions on Automatic Control*, 39(4):797–803, 1994. [49]
- [240] A. Talbot. The evaluation of integrals of products of linear system responses - part I. *Quart. J. Mech. Appl. Math.*, 12(4):488–503, 1959. [51]
- [241] F. De Terán and F. M. Dopico. The solution of the equation $XA + AX^T = 0$ and its application to the theory of orbits. *Linear Algebra Appl.*, 434(1):44–67, 2011. [52]
- [242] F. De Terán and Froilán M. Dopico. Consistency and efficient solution of the Sylvester equation for \star -congruence. *Electronic J. of Linear Algebra*, 22:849–863, 2011. [51]
- [243] Fernando De Terán, Froilán M. Dopico, Nathan Guillery, Daniel Montealegre, and Nicolás Reyes. The solution of the equation $AX + X^*B = 0$. *Lin. Alg. Appl.*, 438(7):2817–2860, 2013. [50]
- [244] Ninoslav Truhar, Zoran Tomljanović, and Ren-Cang Li. Analysis of the solution of the Sylvester equation using low-rank ADI with exact shifts. *Systems & Control Letters*, 59(3-4):248–257, March-April 2010. [13]

- [245] Ninoslav Truhar and Krešimir Veselić. Bounds on the trace of a solution to the Lyapunov equation with a general stable matrix. *Systems & Control Letters*, 56(7-8):493–503, July 2007. [37]
- [246] C.-C. Tsui. A new approach to robust observer design. *Internat. J. Control*, 47:745–751, 1988. [49]
- [247] A. van den Boom, A. Brown, F. Dumortier, A. Geurts, S. Hammarling, R. Kool, M. Vanbegin, P. Van Dooren, and S. Van Huffel. SLICOT, a subroutine library for control and system theory. In *Preprints IFAC Symp. on CADCS*, pages 89–94, Swansea, UK, July 1991. [14, 52]
- [248] P. van Dooren. Structured linear algebra problems in digital signal processing. In *Numerical Linear Algebra, Digital Signal Processing and Parallel Algorithms*, NATO Series F, pages 361–384. Springer, 1991. [5]
- [249] P. van Dooren and M. Verhaegen. Condensed forms for efficient time-invariant Kalman filtering. *SIAM J. Sci. Statist. Comput.*, 9(3):516–530, May 1988. [48]
- [250] Bart Vandereycken and Stefan Vandewalle. A Riemannian optimization approach for computing low-rank solutions of Lyapunov equations. *SIAM J. Matrix Anal. Appl.*, 31(5):2553–2579, 2010. [41]
- [251] Andras Varga. A note on Hammarling’s algorithm for the discrete Lyapunov equation. *Systems & Control Letters*, 15(3):273–275, Sept. 1990. [43]
- [252] D. Vasilyev and J. White. A more reliable reduction algorithm for behavior model extraction. In *Proceedings of the IEEE/ACM International Conference on Computer-Aided Design*, pages 813–820, 2005. [41]
- [253] Yu. O. Vorontsov and Kh. D. Ikramov. A numerical algorithm for solving the matrix equation $AX + X^T B = C$. *Computational Mathematics and Mathematical Physics*, 51(5):691–698, 2011. [50]
- [254] K. w. E. Chu. The solution of the matrix equations $AXB - CXD = E$ and $(YA - DZ, YC - BZ) = (E, F)$. *Lin. Alg. Appl.*, 93:93–105, 1987. [45, 47]
- [255] E. L. Wachspress. Iterative Solution of the Lyapunov Matrix Equations. *Appl. Math. Lett.*, 1(1):87–90, 1988. [35]
- [256] Eugene L. Wachspress. Extended application of Alternating Direction Implicit iteration model Problem theory. *Journal of the Society for Industrial and Applied Mathematics*, 11(4):994–1016, 1963. [8, 15]
- [257] F. Wan. An in-core finite difference method for separable boundary value problems on a rectangle. *Studies in Appl. Math.*, 52:103–113, 1973. [15]
- [258] Qing-Wen Wang, Jian-Hua Sun, and Shang-Zhi Li. Consistency for bi(skew)symmetric solutions to systems of generalized Sylvester equations over a finite central algebra. *Lin. Alg. Appl.*, 353:169–182, 2002. [48]
- [259] David. S. Watkins. *The matrix eigenvalue problem*. SIAM, Philadelphia, 2007. [15]
- [260] H. K. Wimmer. Linear matrix equations, controllability and observability, and the rank of solutions. *SIAM J. Matrix Anal.*, 9(4):570–578, 1988. [12]
- [261] H. K. Wimmer. Explicit solutions of the matrix equation $\sum A^i X D_i = C$. *SIAM J. Matrix Anal. Appl.*, 13:1123–1130, 1992. [12]
- [262] H. K. Wimmer. The generalized Sylvester equation in polynomial matrices. *IEEE Trans. Automat. Control*, 41:1372–1376, 1996. [46]
- [263] H.K. Wimmer. The matrix equation $X - AXB = C$ and an analogue of Roth’s theorem. *Lin. Alg. Appl.*, 109:145–147, 1988. [42]
- [264] H.K. Wimmer and A.D. Ziebur. Remarks on inertia theorems for matrices. *Czechoslovak Math. J.*, 25:556–561, 1975. [42]

- [265] A.-G. Wu, J. Hu, and G.-R. Duan. Solutions to the matrix equation $AX - EXF = BY$. *Computers and Mathematics with Applications*, 58(10):1891–1900, 2009. [47]
- [266] A.-G. Wu, Y. Sun, and G. Feng. Closed-form solution to the non-homogeneous generalised Sylvester matrix equation. *Control Theory & Applications, IET*, 4(10):1914 – 1921, October 2010. [47]
- [267] Ai-Guo Wu, Guang-Ren Duan, and Yu Xue. Kronecker maps and Sylvester-polynomial matrix equations. *IEEE Transactions on Automatic Control*, 52(5), May 2007. [49]
- [268] Ai-Guo Wu, Gang Feng, Guang-Ren Duan, and Wei-Jun Wu. Closed-form solutions to Sylvester-conjugate matrix equations. *Computers and Mathematics with Applications*, 60:95–111, 2010. [51]
- [269] B. Zhou, G.-R. Duan, and Z.-Y. Li. A Stein matrix equation approach for computing coprime matrix fraction description. *Control Theory & Applications, IET*, 3(6):691–700, 2009. [7]
- [270] Bin Zhou and Guang-Ren Duan. Parametric solutions to the generalized discrete Sylvester matrix equation $MXN - X = TY$ and their applications. *IMA Journal of Mathematical Control and Information*, 26(1):59–78, 2009. [47]
- [271] Bin Zhou, Guang-Ren Duan, and Zhao-Yan Li. Gradient based iterative algorithm for solving coupled matrix equations. *Systems & Control Letters*, 58(5):327–333, May 2009. [42, 47, 48]
- [272] Bin Zhou, James Lam, and Guang-Ren Duan. Toward Solution of Matrix Equation $X = Af(X)B + C$. Technical report, arXiv:1211.0346, 2012. [44]
- [273] Yunkai Zhou. *Numerical methods for large scale matrix equations with applications in LTI system model reduction*. PhD thesis, Rice University, Houston, Texas, May 2002. [30, 31]