

A compact rational Krylov method for large-scale rational eigenvalue problems

Froilán M. Dopico*

Javier González-Pizarro†

May 19, 2017

Abstract

In this work, we propose a new method, termed as R-CORK, for the numerical solution of large-scale rational eigenvalue problems, which is based on a linearization and on a compact decomposition of the rational Krylov subspaces corresponding to this linearization. R-CORK is an extension of the compact rational Krylov method (CORK) introduced very recently in [28] to solve a family of non-linear eigenvalue problems that can be expressed and linearized in certain particular ways and which include arbitrary polynomial eigenvalue problems, but not arbitrary rational eigenvalue problems. The R-CORK method exploits the structure of the linearized problem by representing the Krylov vectors in a compact form in order to reduce the cost of storage, resulting in a method with two levels of orthogonalization. The first level of orthogonalization works with vectors of the same size that the original problem, and the second level works with vectors of size much smaller than the original problem. Since vectors of the size of the linearization are never stored or orthogonalized, R-CORK is more efficient from the point of views of memory and orthogonalization than the classical rational Krylov method applied directly to the linearization. Taking into account that the R-CORK method is based on a classical rational Krylov method, to implement implicit restarting is also possible and we show how to do it in a memory efficient way. Finally, some numerical examples are included in order to show that the R-CORK method performs satisfactorily in practice.

Key words. large-scale, linearization, rational eigenvalue problem, rational Krylov method

AMS subject classification. 65F15, 65F50, 15A22

1 Introduction

In this work, we consider the rational eigenvalue problem (REP)

$$R(\lambda)x = 0, \tag{1.1}$$

where $R(\lambda) \in \mathbb{C}(\lambda)^{n \times n}$ is a nonsingular rational matrix, i.e., the entries of $R(\lambda)$ are scalar rational functions in the variable λ with complex coefficients and $\det(R(\lambda)) \not\equiv 0$ is not identically zero, and $x \in \mathbb{C}^n$ is a nonzero vector. More precisely, we consider that $R(\lambda)$ is given as

$$R(\lambda) = P(\lambda) - \sum_{i=1}^k \frac{f_i(\lambda)}{g_i(\lambda)} E_i, \tag{1.2}$$

*Departamento de Matemáticas, Universidad Carlos III de Madrid, Avda. Universidad 30, 28911 Leganés, Spain (dopico@math.uc3m.es). Supported by “Ministerio de Economía, Industria y Competitividad of Spain” and “Fondo Europeo de Desarrollo Regional (FEDER) of EU” through grants MTM2012-32542, MTM2015-65798-P and MTM2015-68805-REDT.

†Departamento de Matemáticas, Universidad Carlos III de Madrid, Avda. Universidad 30, 28911 Leganés, Spain (jagpizar@math.uc3m.es). Supported by “Ministerio de Economía, Industria y Competitividad of Spain” and “Fondo Europeo de Desarrollo Regional (FEDER) of EU” through grants MTM2012-32542 and MTM2015-65798-P and by “Comisión Nacional de Investigación Científica y Tecnológica (CONICYT) de Chile” through grant BCH 72160331.

where $P(\lambda) \in \mathbb{C}[\lambda]^{n \times n}$ is a matrix polynomial of degree d in the variable λ , $f_i(\lambda)$, $g_i(\lambda)$ are coprime scalar polynomials of degrees m_i and n_i , respectively, $m_i < n_i$ and $E_i \in \mathbb{C}^{n \times n}$ are constant matrices for $i = 1, \dots, k$. We emphasize that it is well known that every rational matrix can be written in the form (1.2) [12, 20] (see also [3, Section 2]) and that such form appears naturally in many applications [26].

The REP has attracted considerable interest in recent years since it arises in different applications in some fields such as vibration of fluid-solid structures [29], optimization of acoustic emissions of high speed trains [15], free vibration of plates with elastically attached masses [23], free vibrations of a structure with a viscoelastic constitutive relation describing the behavior of a material [18, 19], and electronic structure calculations of quantum dots [11, 30].

A first idea to solve REPs is a brute-force approach, since one can multiply by $\prod_{i=1}^k g_i(\lambda)$ to turn the rational matrix (1.2) into a matrix polynomial of degree $d + n_1 + \dots + n_k$. The common approach to solve a polynomial eigenvalue problem (PEP) is via linearization (see, for instance, [7, 16, 18]), this is, by transforming the PEP into a generalized eigenvalue problem (GEP) and then applying a well-established algorithm to this GEP, as for instance the QZ algorithm in the case of dense medium sized problems [9] or some Krylov subspace method for large-scale problems. However, this brute-force approach it is only useful when $n_1 + n_2 + \dots + n_k$ is small compared with d . So, if k or some n_i are big, then the degree of the matrix polynomial associated to the problem is also big, and this makes the size of the linearization too large, which is impractical for medium to large-scale problems. This drawback has motivated the idea of linearizing directly the REP [26]. The linearization for $R(\lambda)$ in (1.2) constructed in [26] has a size much smaller than the size of the linearization obtained by the brute-force approach. Nonetheless the increase of the size of the problem is still considerable, so for large-scale rational eigenvalue problems, a direct application of this approach, i.e., without taking into account the structure of the linearization, is also impractical. This idea of taking advantage of the structure of the linearization for solving large-scale REPs is closely connected to the intense research effort developed in the last years by different authors for solving large-scale PEPs via linearizations and that is briefly discussed in the next paragraph.

Several methods have been developed to solve large-scale PEPs numerically by applying Krylov methods to the associated GEPs obtained through linearizations. In this approach, the key issues to be solved for using Krylov methods for large-scale PEPs are the increase of the memory cost and the increase of the orthogonalization cost at each step, as a consequence of the increase of the size of the linearization with respect to the size of the original problem. In order to reduce these costs, different representations of the Krylov vectors of the linearizations have been developed. First, the second order Arnoldi method (SOAR) [4] and the quadratic Arnoldi method (Q-Arnoldi) [17] were developed to solve quadratic eigenvalue problems (QEP), introducing a new representation of the Krylov vectors. However, both methods are potentially unstable as a consequence of performing implicitly the orthogonalization. To cure this instability, the two-level orthogonal Arnoldi process (TOAR) [27, 14] for QEP proposed a different compact representation for the Krylov vectors of the linearization and, combining this representation with the linearization and the Arnoldi recurrence relation, resulted in a memory saving and numerically stable method. Extending the ideas of a compact representation of the Krylov vectors and of the two levels of orthogonalization from polynomials of degree 2 (TOAR) to polynomials of any degree, the authors of [13] developed a memory-efficient and stable Arnoldi process for linearizations of matrix polynomials expressed in the Chebyshev basis. In 2015, the compact rational Krylov method (CORK) for nonlinear eigenvalue problems (NLEP) was introduced in [28]. CORK considers particular NLEPs that can be expressed and linearized in certain ways, which are solved by applying a compact rational Krylov method to such linearizations. A key feature of the CORK method is that it works for many kinds of linearizations involving a Kronecker structure, as the Frobenius companion form or linearizations of matrix polynomials in different bases (as Newton or Chebyshev, among others [2]). CORK reduces

both the costs of memory and orthogonalization by using a generalization of the compact Arnoldi representation of the Krylov vectors of the linearizations used in TOAR [27, 14], and gets stability through two levels of orthogonalization as in TOAR.

In this paper, we develop a rational Krylov method that works on the linearization of REPs introduced in [26] to solve large-scale and sparse $n \times n$ REPs. To this aim, we introduce a compact rational Krylov method for REPs (R-CORK). In the spirit of TOAR and CORK, we will work with two levels of orthogonalization, and, as in CORK, we adapt the classical rational Krylov method [21, 22, 28] on the linearization to a compact representation of the Krylov vectors and to the two levels of orthogonalization. We can perform the shift-and-invert step by solving linear systems of size n . To this purpose, the linearization introduced in [26] is preprocessed in a convenient way and, then, an ULP decomposition is used. This decomposition is similar to the one employed in [28] directly on the linearizations of the NLEPs considered there. Once this step is performed, we start with the two levels of orthogonalization. The first level involves an orthogonalization process with vectors of size n and in the second level of orthogonalization we work with vectors of size much smaller than n , so this level is cheap compared with the first level. As a result, we develop a stable method that allows us to reduce the orthogonalization cost and the memory cost by exploiting the structure of the matrix pencil that linearizes the REP and using the rational Krylov recurrence relation.

The rest of the paper is organized as follows. Section 2 introduces some preliminary concepts: a summarized background on polynomial and rational eigenvalue problems, the classical rational Krylov method for the generalized eigenvalue problem, and the CORK method particularized to polynomial eigenvalue problems. Section 3 proposes the compact rational Krylov decomposition that we use to develop the R-CORK method and presents the detailed algorithm with the two levels of orthogonalization. Section 4 discusses the implementation of implicit restarting for the R-CORK method. Section 5 presents numerical examples which show that the R-CORK method works satisfactorily in practice and, finally, the main conclusions and some lines of future research are discussed in Section 6.

Notation. We denote vectors by lowercase characters, u , and matrices by capital characters, A . Block vectors and block matrices are denoted by bold face fonts, \mathbf{u} , and \mathbf{A} , respectively, and the i -th block of \mathbf{u} is represented by $u^{(i)}$. The conjugate transpose of A is denoted as A^* . The $i \times j$ matrix with the main diagonal entries equal to 1 and the rest of entries equal to zero is represented by $I_{i \times j}$. In the particular case of $i = j$ this matrix is the identity matrix and is denoted by I_i . The vector e_j represents the canonical vector associated to the j -th column of the identity matrix and $0_{i \times j}$ represents the zero matrix of size $i \times j$, which in the particular case $i = j$ is denoted simply by 0_j . The matrix U_j represents a matrix with j columns and u_i represents the i -th column of U_j . The rational Krylov subspace of order m associated with the matrices A and $B \in \mathbb{C}^{n \times n}$, the initial vector $u_1 \in \mathbb{C}^n$ and the shifts $\theta_1, \theta_2, \dots, \theta_{m-1} \in \mathbb{C}$ is denoted by

$$\mathcal{K}_m(A, B, u_1, \theta_1, \dots, \theta_{m-1}) = \text{span}\{u_1, (A - \theta_1 B)^{-1} B u_1, (A - \theta_2 B)^{-1} B u_2, \dots, (A - \theta_{m-1} B)^{-1} B u_{m-1}\} \quad (1.3)$$

where $u_{i+1} = (A - \theta_i B)^{-1} B u_i$, $i = 1, \dots, m - 1$. We omit subscripts when the dimensions of the matrices are clear from the context. The norm $\|\cdot\|_2$ represents the 2-norm and $\|\cdot\|_F$ the Frobenius norm [10, Ch. 5]. The Kronecker product of two matrices is denoted by $A \otimes B$. The set of $n \times n$ rational matrices is denoted by $\mathbb{C}(\lambda)^{n \times n}$ and the set of $n \times n$ polynomial matrices (or, equivalently, matrix polynomials) is denoted by $\mathbb{C}[\lambda]^{n \times n}$.

2 Preliminaries

2.1 Basics on polynomial eigenvalue problems and linearizations

The classical approach to solve a regular PEP

$$P(\lambda)x = 0, \quad (2.1)$$

where $P(\lambda) = \sum_{i=0}^d \lambda^i P_i$ with $P_i \in \mathbb{C}^{n \times n}$ and $\det(P(\lambda)) \not\equiv 0$ is via linearization. In this process, the matrix polynomials are mapped into matrix pencils with the same eigenvalues and multiplicities [8, 16]. More precisely, a pencil $L(\lambda) = \mathbf{A} - \lambda \mathbf{B}$ is called a linearization of $P(\lambda)$ if there exist unimodular matrix polynomials¹ $E_1(\lambda), E_2(\lambda)$ such that

$$\begin{bmatrix} P(\lambda) & 0 \\ 0 & I_{(d-1)n} \end{bmatrix} = E_1(\lambda)(\mathbf{A} - \lambda \mathbf{B})E_2(\lambda).$$

Some linearizations of matrix polynomials of degree d and size $n \times n$, very useful in practice, are of the form as the pencils in Definition 2.1.

Definition 2.1. [28, Definition 2.2] Let $P(\lambda) \in \mathbb{C}[\lambda]^{n \times n}$ be a regular matrix polynomial, i.e., $\det(P(\lambda))$ does not vanish identically, of degree $d \geq 2$ and size $n \times n$. A $dn \times dn$ matrix pencil $L(\lambda)$ of the form

$$L(\lambda) = \mathbf{A} - \lambda \mathbf{B}, \quad (2.2)$$

where

$$\mathbf{A} = \begin{bmatrix} A_0 & A_1 & \cdots & A_{d-1} \\ \hline & M \otimes I_n & & \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} B_0 & B_1 & \cdots & B_{d-1} \\ \hline & N \otimes I_n & & \end{bmatrix} \quad (2.3)$$

and $A_i, B_i \in \mathbb{C}^{n \times n}$, $i = 0, 1, \dots, d-1$, and $M, N \in \mathbb{C}^{(d-1) \times d}$, is called a structured linearization pencil of $P(\lambda)$ if the following conditions hold

- 1) $L(\lambda)$ is a linearization of $P(\lambda)$,
- 2) $M - \lambda N$ has rank $d-1$ for all $\lambda \in \mathbb{C}$, and
- 3) $(\mathbf{A} - \lambda \mathbf{B})(f(\lambda) \otimes I_n) = e_1 \otimes P(\lambda)$ for some polynomial function $f : \mathbb{C} \rightarrow \mathbb{C}[\lambda]^d$, $f(\lambda) \neq 0$ for all $\lambda \in \mathbb{C}$, where $e_1 \in \mathbb{C}^d$ is the first vector of the canonical basis of \mathbb{C}^d .

The matrices A_i and B_i that appear in the first block rows in (2.3) are related to the matrix polynomial $P(\lambda)$ and the matrices M and N correspond to the linear relations between the basis functions $f_i(\lambda)$, where $f(\lambda) := [f_1(\lambda), \dots, f_d(\lambda)]^T$, used in the representation of the matrix polynomial. The interested reader can find some examples in [28]. The identity $(\mathbf{A} - \lambda \mathbf{B})(f(\lambda) \otimes I_n) = e_1 \otimes P(\lambda)$ generalizes the identity used in [16] to define certain vector spaces of linearizations of matrix polynomials.

An important property of structured linearization pencils is that their eigenvectors are closely related to the eigenvectors of the matrix polynomial as we can see in Theorem 2.2.

Theorem 2.2. [28, Corollary 2.4] Let $L(\lambda)$ be a structured linearization pencil of $P(\lambda)$ as in Definition 2.1 and let (λ_*, \mathbf{x}) be an eigenpair of $L(\lambda)$. Then, the eigenvector \mathbf{x} has the following structure

$$\mathbf{x} = f(\lambda_*) \otimes x,$$

where $x \in \mathbb{C}^n$ is an eigenvector of $P(\lambda)$ corresponding to λ_* .

¹Unimodular matrix polynomials are matrix polynomials whose determinant is a nonzero constant, i.e., it does not depend of λ . Most of the linearizations considered in this work are in fact strong linearizations [7, 16], so, they preserve also the eigenvalues at infinity of $P(\lambda)$, and their multiplicities, if they are present. Nevertheless in this work we do not intend to compute infinite eigenvalues since their existence is not generic, and so we do not need to use the concept of strong linearization.

The block ULP decomposition in Theorem 2.3 for structured linearization pencils of matrix polynomials is important for the CORK method introduced in [28] because it allows to perform the shift-and-invert step in CORK efficiently. We will use also a decomposition of this type to perform the shift-and-invert step in the R-CORK method developed in Section 3.

Theorem 2.3. [28, Theorem 2.3] *Let \mathbf{A} and \mathbf{B} be defined by (2.3). Then, for every $\mu \in \mathbb{C}$ there exists a permutation matrix $\mathcal{P} \in \mathbb{C}^{d \times d}$ such that the matrix $(M_1 - \mu N_1) \in \mathbb{C}^{(d-1) \times (d-1)}$ is invertible with*

$$M =: [m_0 \quad M_1]\mathcal{P}, \quad N =: [n_0 \quad N_1]\mathcal{P}.$$

Moreover, the matrix $L(\mu)$, i.e., the pencil $L(\lambda)$ in (2.2) evaluated in μ , can be factorized as follows

$$L(\mu) = \mathbf{A} - \mu\mathbf{B} = \mathcal{U}(\mu)\mathcal{L}(\mu)(\mathcal{P} \otimes I_n),$$

where

$$\begin{aligned} \mathcal{L}(\mu) &= \begin{bmatrix} P(\mu) & 0 \\ (m_0 - \mu n_0) \otimes I_n & (M_1 - \mu N_1) \otimes I_n \end{bmatrix}, \\ \mathcal{U}(\mu) &= \begin{bmatrix} \alpha^{-1} I_n & (\bar{\mathbf{A}}_1 - \mu \bar{\mathbf{B}}_1)((M_1 - \mu N_1)^{-1} \otimes I_n) \\ 0 & I_{(d-1)n} \end{bmatrix}, \end{aligned}$$

with the scalar $\alpha = e_1^T \mathcal{P} f(\mu) \neq 0$ and

$$\begin{aligned} [A_0 \quad A_1 \quad \cdots \quad A_{d-1}] &=: [\bar{A}_0 \quad \bar{\mathbf{A}}_1](\mathcal{P} \otimes I_n), \\ [B_0 \quad B_1 \quad \cdots \quad B_{d-1}] &=: [\bar{B}_0 \quad \bar{\mathbf{B}}_1](\mathcal{P} \otimes I_n). \end{aligned}$$

2.2 A linearization for rational eigenvalue problems

In this subsection, we present some results and notations related to the REP. Interested readers can find more information in the summaries presented in [1, Sections 1 & 2] and [3, Section 2], as well as in the classical references [12, 20].

In this work, we assume that the rational matrix $R(\lambda)$ in (1.2) is regular, this means, $\det(R(\lambda)) \not\equiv 0$. With a slight lack of rigor, we can say that if the matrices E_i in (1.2) are linearly independent, then the roots of the denominators $g_i(\lambda)$ are the poles of $R(\lambda)$ and that $R(\lambda)$ is not defined in these poles. A scalar $\lambda \in \mathbb{C}$ which is not a pole is called an eigenvalue of $R(\lambda)$ if $\det(R(\lambda)) = 0$, and a nonzero vector $x \in \mathbb{C}^n$ is called an eigenvector of $R(\lambda)$ associated to the eigenvalue λ if the condition (1.1) holds. The pair (λ, x) constitutes an eigenpair of $R(\lambda)$ and our goal is to compute a subset of such eigenpairs.

We express the matrix polynomial $P(\lambda)$ of degree d in (1.2) as follows

$$P(\lambda) = \lambda^d P_d + \lambda^{d-1} P_{d-1} + \cdots + \lambda P_1 + P_0, \quad (2.4)$$

where $P_i \in \mathbb{C}^{n \times n}$ for $i = 0, \dots, d$. From now on, we assume the generic condition that the leading coefficient matrix P_d is nonsingular in (2.4). As explained in the introduction, we assume that $f_i(\lambda)$ and $g_i(\lambda)$ in (1.2) are coprime, this is, they do not have common factors, and that the rational functions $\frac{f_i(\lambda)}{g_i(\lambda)}$ are strictly proper, this is, the degree, m_i , of $f_i(\lambda)$ is smaller than the degree, n_i , of $g_i(\lambda)$. Under these assumptions, in [26], Su and Bai proposed a linearization to solve the rational eigenvalue problem. With this aim, they first showed that one can find matrices E, F of size $n \times s$, and matrices C, D of size $s \times s$, with $s = r_1 n_1 + r_2 n_2 + \cdots + r_k n_k$ with $r_i = \text{rank}(E_i)$ in (1.2), such that

$$R(\lambda) = P(\lambda) - E(C - \lambda D)^{-1} F^T. \quad (2.5)$$

In fact, it is a classical result (much older than [26]) that any rational matrix can be written as in (2.5) by expressing $R(\lambda)$ as the sum of its unique polynomial and strictly proper parts, and,

then, constructing a state-space realization of the strictly proper part [20] (see also [3, Section 2]). However, we emphasize that, as far as we know, [26] is the first reference available in the literature that uses (2.5) with the purpose of computing the eigenvalues of a REP, as well as that [26] is the first reference that points out that the representation (2.5) is immediately available from the data in many practical REPs without any computational cost.

Once the representation (2.5) for the REP is available, the authors of [26] linearized the REP $R(\lambda)x = 0$ as follows:

$$(\mathcal{A} - \lambda\mathcal{B})\mathbf{z} = 0, \quad (2.6)$$

where

$$\mathcal{A} = \left[\begin{array}{cccc|c} P_{d-1} & P_{d-2} & \cdots & P_0 & E \\ -I_n & 0 & \cdots & 0 & \\ & \ddots & \ddots & \vdots & \\ & & -I_n & 0 & \\ \hline & & & F^T & C \end{array} \right], \quad \mathcal{B} = - \left[\begin{array}{ccc|c} P_d & & & \\ & I_n & & \\ & & \ddots & \\ & & & I_n \\ \hline & & & -D \end{array} \right] \quad (2.7)$$

and

$$\mathbf{z} = \left[\begin{array}{c} \lambda^{d-1}x \\ \lambda^{d-2}x \\ \vdots \\ x \\ \hline y \end{array} \right], \quad \text{with } y = -(C - \lambda D)^{-1} F^T x. \quad (2.8)$$

Denoting by \mathbf{A} and \mathbf{B} the upper left $nd \times nd$ submatrices of \mathcal{A} and \mathcal{B} , we can write $\mathcal{A} - \lambda\mathcal{B}$ as follows

$$\mathcal{A} - \lambda\mathcal{B} = \left[\begin{array}{c|c} \mathbf{A} - \lambda\mathbf{B} & e_1 \otimes E \\ \hline e_d^T \otimes F^T & C - \lambda D \end{array} \right], \quad (2.9)$$

where e_1 and e_d are the first and the last columns of I_d , respectively. Observe that $\mathbf{A} - \lambda\mathbf{B}$ is the famous first Frobenius companion linearization of the matrix polynomial $P(\lambda)$ in (2.4) [8].

As mentioned above, it is important to remark that in many applications of REPs [18, 26, 30], the first step in the process above, i.e., to construct the representation (2.5), does not involve any computational effort, since the matrices $P_0, P_1, \dots, P_d, E, C, D$, and F can be obtained directly from the data. As a consequence the linearization $\mathcal{A} - \lambda\mathcal{B}$ above can be constructed also without any computational effort and all the computational effort is attached to the solution of the GEP (2.6). Another important remark that has a deep computational impact is that in many applications of REPs [18, 26, 30] the size $s \times s$ of the matrices C and D is much smaller than the size $n \times n$ of the original REP, i.e., $s \ll n$ or, in plain words, the rank of the strictly proper part of $R(\lambda)$ is much smaller than the size of $R(\lambda)$. Therefore, if $s \ll n$, then the size $(nd + s) \times (nd + s)$ of the linearization $\mathcal{A} - \lambda\mathcal{B}$ is approximately equal to the size $(nd) \times (nd)$ of the linearization $\mathbf{A} - \lambda\mathbf{B}$ of the matrix polynomial $P(\lambda)$, and the costs of solving the GEPs $\mathcal{A} - \lambda\mathcal{B}$ and $\mathbf{A} - \lambda\mathbf{B}$ are expected to be similar. Thus, it is not surprising that the R-CORK algorithm developed in Section 3 for large-scale REPs is particularly efficient in terms of storage and orthogonalization costs when $s \ll n$. However, we emphasize in this context that R-CORK also improves significantly these costs when $s \approx n$ (and $d \geq 2$) with respect to a direct application of large-scale eigensolvers to $\mathcal{A} - \lambda\mathcal{B}$. We will make often comments about the advantages of considering $s \ll n$ throughout the paper.

A formal definition of linearization of a rational matrix can be found in [1] and another one which includes the concept of strong linearization in [3]. In fact, it is proved in [3] that $\mathcal{A} - \lambda\mathcal{B}$ in (2.6) is a strong linearization of $R(\lambda)$ in (2.5) whenever $-E(C - \lambda D)^{-1}F^T$ is a *minimal* state-space realization [20] of the strictly proper part of $R(\lambda)$ [3, Section 8]. We emphasize that the requirement that $-E(C - \lambda D)^{-1}F^T$ is a minimal state-space realization is very mild [3, Section 8] and that is fully necessary to guarantee that for every eigenvalue λ of $R(\lambda)$ the matrix $C - \lambda D$ is nonsingular

[3, Example 3.2]. In the rest of the paper we implicitly assume that $E(C - \lambda D)^{-1}F^T$ is a minimal realization, although the only result we will use explicitly is Theorem 2.4, which remains valid even when $E(C - \lambda D)^{-1}F^T$ is not minimal. The subtle point is that Theorem 2.4 assumes that λ is a number such that the matrix $(C - \lambda D)$ is invertible, but if $E(C - \lambda D)^{-1}F^T$ is not minimal then there may be eigenvalues of $R(\lambda)$ that do not satisfy such assumption.

Theorem 2.4. [26, Theorem 3.1] *Let $\lambda \in \mathbb{C}$ be such that $\det(C - \lambda D) \neq 0$. Then the following statements hold:*

- (a) *If λ is an eigenvalue of the REP (2.5), then it is an eigenvalue of the GEP (2.6).*
- (b) *Let λ be an eigenvalue of the GEP (2.6) and $z = [z_1^T, z_2^T, \dots, z_d^T, y^T]^T$ be a corresponding eigenvector, where z_i are vectors of length n for $i = 1, 2, \dots, d$, and y is a vector of length s . Then $z_d \neq 0$ and $R(\lambda)z_d = 0$, namely, λ is an eigenvalue of the REP (2.5) and z_d is a corresponding eigenvector. Moreover, the algebraic and geometric multiplicities of λ for the REP (2.5) and GEP (2.6) are the same.*

Part (b) of Theorem 2.4 is the key result that allows us to get the eigenvalues and eigenvectors of $R(\lambda)$ from those of the GEP $\mathcal{A} - \lambda\mathcal{B}$ in (2.6). In fact, observe that Theorem 2.4-(b) can be improved if $\lambda \neq 0$, since in this case, according to (2.8), every z_i is an eigenvector of $R(\lambda)$. Therefore, we have d degrees of freedom for the recovery of the eigenvector of $R(\lambda)$. The most sensible option from the point of view of rounding errors is to choose the z_i with largest 2-norm, that is, z_1 when $|\lambda| > 1$ and z_d when $|\lambda| \leq 1$.

The final comment of this section is that in contrast to the CORK method developed in [28] for PEPs, which is valid for many linearizations, the rational CORK method, R-CORK, introduced in this manuscript uses only the linearization $\mathcal{A} - \lambda\mathcal{B}$ in (2.6). The reason of this restriction is that the theory of linearizations of REPs is far less developed than the theory of linearizations of PEPs. Thus, although many linearizations of REPs have been introduced very recently in [1, 3], their properties are not yet fully understood.

2.3 The classical rational Krylov method for generalized eigenvalue problems

We revise in this subsection the rational Krylov method for GEPs since the algorithm R-CORK presented in this paper is based on this method. The rational Krylov method [21, 22] is a generalization for computing eigenvalues of matrices and of matrix pencils of the shift-and-invert Arnoldi method. The main differences between these methods are basically two: in rational Krylov methods we can change the shift θ_j at each iteration instead of fixing the shift as in the shift-and-invert Arnoldi method. Also, the information of the approximate eigenvalues is contained in two upper Hessenberg matrices \underline{H}_j and \underline{K}_j instead of in only one matrix. In Algorithm 1 we present a basic pseudocode of the rational Krylov method that summarizes its main steps and guides the developments in the rest of this subsection, which are a very brief sketch of the rational Krylov method. The reader can find more details in [21, 22, 28]. This method produces an orthonormal basis for the subspace $\mathcal{K}_{m+1}(\mathcal{A}, \mathcal{B}, \mathbf{u}_1, \theta_1, \dots, \theta_m)$. By using the equalities for $\hat{\mathbf{u}}$ and $\tilde{\mathbf{u}}$ from steps 2 and 4 in Algorithm 1 we obtain at the i -th iteration:

$$(\mathcal{A} - \theta_i \mathcal{B})^{-1} \mathcal{B} \mathbf{u}_i = \mathbf{U}_{i+1} \underline{h}_i,$$

with $\underline{h}_i = [h_i^* \quad h_{i+1,i}^*]^*$. After j steps of the rational Krylov method, we obtain the classic rational Krylov recurrence relation [22]:

$$\mathcal{A} \mathbf{U}_{j+1} \underline{H}_j = \mathcal{B} \mathbf{U}_{j+1} \underline{K}_j, \quad (2.10)$$

where $\underline{H}_j, \underline{K}_j \in \mathbb{C}^{(j+1) \times j}$ are upper Hessenberg matrices and

$$\underline{K}_j = \underline{H}_j \text{diag}(\theta_1, \theta_2, \dots, \theta_j) + I_{(j+1) \times j}. \quad (2.11)$$

Algorithm 1 Rational Krylov method

Input: \mathcal{A} and \mathcal{B} square matrices and an initial vector \mathbf{u}_1 with $\|\mathbf{u}_1\|_2 = 1$.

Output: The matrix \mathbf{U}_{m+1} whose columns are an orthonormal basis of $\mathcal{K}_{m+1}(\mathcal{A}, \mathcal{B}, \mathbf{u}_1, \theta_1, \dots, \theta_m)$, and the Ritz pairs (λ, \mathbf{x}) of $\mathcal{A} - \lambda\mathcal{B}$, corresponding to the rational Krylov subspace $\mathcal{K}_m(\mathcal{A}, \mathcal{B}, \mathbf{u}_1, \theta_1, \dots, \theta_{m-1})$.

Initialize $\mathbf{U}_1 = [\mathbf{u}_1]$.

for $j = 1, 2, \dots, m$ **do**

1. Choose the shift θ_j .

2. $\hat{\mathbf{u}} = (\mathcal{A} - \theta_j\mathcal{B})^{-1}\mathcal{B}\mathbf{u}_j$.

3. $h_j = \mathbf{U}_j^*\hat{\mathbf{u}}$.

4. $\tilde{\mathbf{u}} = \hat{\mathbf{u}} - \mathbf{U}_j h_j$.

5. Compute the new vector $\mathbf{u}_{j+1} = \tilde{\mathbf{u}}/h_{j+1,j}$ with $h_{j+1,j} = \|\tilde{\mathbf{u}}\|_2$.

6. Update $\mathbf{U}_{j+1} = [\mathbf{U}_j \quad \mathbf{u}_{j+1}]$.

7. Compute the eigenpairs (λ_i, t_i) of (2.12) and test for convergence.

end for

8. Compute the eigenvectors $\mathbf{x}_i = \mathbf{U}_{j+1}\underline{H}_j t_i$, $i = 1, \dots, j$.

For simplicity, we assume that breakdowns do not occur in the rational Krylov method, this is, $h_{i+1,i} \neq 0$ for all $i = 1, \dots, m$, and in this case the upper Hessenberg matrix \underline{H}_j is unreduced. We can approximate in each iteration of Algorithm 1 the corresponding j eigenvalues and eigenvectors of the pencil $\mathcal{A} - \lambda\mathcal{B}$ by solving the small generalized eigenvalue problem:

$$K_j t_i = \lambda_i H_j t_i, \quad t_i \neq 0, \quad (2.12)$$

where H_j and K_j are the $j \times j$ upper Hessenberg matrices obtained by removing the last rows of \underline{H}_j and \underline{K}_j , respectively. Then, we call $(\lambda_i, \mathbf{x}_i = \mathbf{U}_{j+1}\underline{H}_j t_i)$ a Ritz pair of $(\mathcal{A}, \mathcal{B})$. We emphasize that the approximate eigenvectors \mathbf{x}_i are not computed in each iteration, since this would be very expensive, and that the test for convergence in step 7 of Algorithm 1 can be performed in an inexpensive way by using only the small vectors t_i , as it is done in most Krylov methods.

2.4 The CORK method for polynomial eigenvalue problems

Van Beeumen, Meerbergen, and Michiels in [28] proposed a method based on a compact rational Krylov decomposition, extending the two levels of orthogonalization idea of TOAR from the quadratic eigenvalue problem [27, 14] to arbitrary degree polynomial eigenvalue problems and to other NLEPs, including many other linearizations apart from the Frobenius one used in [27, 14], and using the rational Krylov method instead of the Arnoldi method. This method was baptized as CORK in [28] and for simplicity we described it particularized to PEPs of degree d . The key idea in [28] is to apply the rational Krylov method in Algorithm 1 to a structured linearization pencil of a matrix polynomial $P(\lambda)$ of degree d (recall Definition 2.1) taking into account that the special structure of these pencils imposes a special structure on the bases of the corresponding rational Krylov subspaces. By using this structure, the authors of [28] reduced both the memory cost and the orthogonalization cost of the classical rational Krylov method applied to an arbitrary pencil of the same size. Considering the matrices \mathbf{A} and \mathbf{B} in (2.3) and the rational Krylov recurrence relation (2.10) for \mathbf{A} and \mathbf{B} , the authors of [28] partitioned conformably the matrix \mathbf{U}_{j+1} as follows

$$\mathbf{U}_{j+1} = [\mathbf{U}_j \quad \mathbf{u}_{j+1}] = \begin{bmatrix} U_j^{(1)} & u_{j+1}^{(1)} \\ U_j^{(2)} & u_{j+1}^{(2)} \\ \vdots & \vdots \\ U_j^{(d)} & u_{j+1}^{(d)} \end{bmatrix},$$

and then, they constructed a matrix $Q_j \in \mathbb{C}^{n \times r_j}$ with orthonormal columns such that

$$\text{span}\{Q_j\} = \text{span}\{U_j^{(1)}, U_j^{(2)}, \dots, U_j^{(d)}\} \quad (2.13)$$

and $\text{rank}(Q_j) = r_j$. By using the matrix Q_j , the blocks $U_j^{(i)}$ for $i = 1, 2, \dots, d$ can be represented as follows

$$U_j^{(i)} = Q_j R_j^{(i)}, \quad i = 1, 2, \dots, d,$$

for some matrices $R_j^{(i)} \in \mathbb{C}^{r_j \times j}$. Then,

$$\mathbf{U}_j = \begin{bmatrix} Q_j R_j^{(1)} \\ Q_j R_j^{(2)} \\ \vdots \\ Q_j R_j^{(d)} \end{bmatrix} = \begin{bmatrix} Q_j & & & \\ & Q_j & & \\ & & \ddots & \\ & & & Q_j \end{bmatrix} \begin{bmatrix} R_j^{(1)} \\ R_j^{(2)} \\ \vdots \\ R_j^{(d)} \end{bmatrix} = (I_d \otimes Q_j) \mathbf{R}_j, \quad (2.14)$$

where

$$\mathbf{R}_j := \begin{bmatrix} R_j^{(1)} \\ R_j^{(2)} \\ \vdots \\ R_j^{(d)} \end{bmatrix}.$$

By using this representation, the rational Krylov recurrence relation (2.10) can be written as follows [28, eq. (4.3)]

$$\mathbf{A}(I_d \otimes Q_{j+1}) \mathbf{R}_{j+1} \underline{H}_j = \mathbf{B}(I_d \otimes Q_{j+1}) \mathbf{R}_{j+1} \underline{K}_j.$$

Observe that \mathbf{U}_j has ndj entries while the representation in (2.14) involves $(n + jd)r_j$ parameters. Therefore, taking into account that in the solution of large-scale PEPs the dimension j of the rational Krylov subspaces is much smaller than the dimension n of the problem and that the degree d of applied PEPs is a low number (for sure smaller than 30, see [13], and often much smaller than 30 [5]), we get that $jd \ll n$ and that the representation (2.14) of \mathbf{U}_j stores approximately nr_j numbers. The fundamental reason why the representation of \mathbf{U}_j in (2.14) is of interest and is indeed compact is because r_j is considerably much smaller than jd for the matrices \mathbf{A} and \mathbf{B} in (2.3). More precisely, the following result is proved in [28].

Theorem 2.5. [28, Theorems 4.4 and 4.5] *Let Q_j be defined as in (2.13). Then*

$$\text{span}\{Q_{j+1}\} = \text{span}\{Q_j, u_{j+1}^{(p)}\}, \quad (2.15)$$

where $u_{j+1}^{(p)}$ represents the block of the vector \mathbf{u}_{j+1} in a certain p -th position determined in [28]. Also,

$$r_j < j + d. \quad (2.16)$$

Note that Theorem 2.5 shows that Q_j can be expanded to Q_{j+1} by orthogonalizing only one vector of size n at each iteration. Also, \mathbf{R}_{j+1} can be expanded in an easy way, if $u_{j+1}^{(p)} \notin \text{span}\{Q_j\}$ then the blocks $R_{j+1}^{(i)}$, $i = 1, \dots, d$, can be written as

$$R_{j+1}^{(i)} = \left[\begin{array}{c|c} R_j^{(i)} & \\ \hline 0_{1 \times j} & r_{j+1}^{(i)} \end{array} \right], \quad i = 1, \dots, d,$$

and, if $u_{j+1}^{(p)} \in \text{span}\{Q_j\}$, then $R_{j+1}^{(i)} = \left[R_j^{(i)} \quad r_{j+1}^{(i)} \right]$, $i = 1, \dots, d$. Based on these ideas, the authors of [28] developed CORK, splitting the method into two levels of orthogonalization: the first level

is to expand Q_j into Q_{j+1} and the second level is to expand \mathbf{R}_j into \mathbf{R}_{j+1} . We can see a basic pseudocode for the CORK method in Algorithm 2, whose complete explanation can be found in [28]. For simplicity, we assume that breakdown does not occur in Algorithm 2, i.e., $h_{j+1,j} \neq 0$ for all j .

Algorithm 2 Compact rational Krylov method (CORK)

Input: $Q_1 \in \mathbb{C}^{n \times r_1}$ and $\mathbf{R}_1 \in \mathbb{C}^{dr_1 \times 1}$ with $Q_1^* Q_1 = I_{r_1}$ and $\mathbf{R}_1^* \mathbf{R}_1 = 1$, where $r_1 \leq d$.

Output: Approximate eigenpairs (λ, \mathbf{x}) associated to $\mathbf{A} - \lambda \mathbf{B}$, with \mathbf{A} , \mathbf{B} as in (2.3).

for $j = 1, 2, \dots$ **do**

1. Choose shift θ_j .

First level of orthogonalization:

2. Compute $\hat{u}^{(p)}$ by using the ULP decomposition in Theorem 2.3 with $\mu = \theta_j$ (see [28] for details).

3. Orthogonalize: $\tilde{q} = \hat{u}^{(p)} - Q_j Q_j^* \hat{u}^{(p)}$.

4. If $\tilde{q} \neq 0$ then compute next vector: $q_{j+1} = \tilde{q} / \|\tilde{q}\|_2$ and $Q_{j+1} = [Q_j \ q_{j+1}]$. Otherwise $Q_{j+1} = Q_j$.

Second level of orthogonalization:

5. If $r_{j+1} > r_j$ then update matrices: $R_j^{(i)} = \begin{bmatrix} R_j^{(i)} \\ 0_{1 \times j} \end{bmatrix}$ for $i = 1, \dots, d$.

6. Compute: $\hat{\mathbf{r}}$ by using the ULP decomposition in Theorem 2.3 (see [28] for details).

7. Compute: $\tilde{\mathbf{r}} = \hat{\mathbf{r}} - \mathbf{R}_j h_j$, where $h_j = \mathbf{R}_j^* \hat{\mathbf{r}}$.

8. Next vector: $\mathbf{r}_{j+1} = \tilde{\mathbf{r}} / h_{j+1,j}$, where $h_{j+1,j} = \|\tilde{\mathbf{r}}\|_2$ and $\mathbf{R}_{j+1} = [\mathbf{R}_j \ \mathbf{r}_{j+1}]$.

9. Compute eigenpairs: (λ_i, t_i) of (2.12) and test for convergence.

end for

10. Compute eigenvectors: $\mathbf{x}_i = (I_d \otimes Q_{j+1}) \mathbf{R}_{j+1} \underline{H}_j t_i$.

From the discussion above, it is clear that CORK reduces significantly the storage requirements with respect to a direct application of the rational Krylov method to the $(nd) \times (nd)$ GEP $\mathbf{A} - \lambda \mathbf{B}$, since essentially CORK represents \mathbf{U}_j in terms of $n(j+d)$ parameters and, in addition, $n(j+d) \approx nj$ for moderate values of d . Therefore, the memory cost of CORK is approximately the cost of any Krylov method applied to an $n \times n$ GEP. Moreover, it can be seen in [28, Section 5.4] that the orthogonalization cost of CORK is essentially independent of d for moderate values of d , and, so, much lower than the orthogonalization cost of a direct application of rational Krylov to $\mathbf{A} - \lambda \mathbf{B}$. With respect to the comparison of the costs of the shift-and-invert steps in CORK (included in step 2 of Algorithm 2) and in rational Krylov (step 2 in Algorithm 1), we can say that in CORK the particular structure of the pencil $\mathbf{A} - \lambda \mathbf{B}$ together with the fact that only one block of the vector $\hat{\mathbf{u}}$ is needed allow us to perform this step very efficiently by essentially solving just one ‘‘difficult’’ $n \times n$ linear system (see [28, Algorithm 2]). In contrast, in rational Krylov the whole vector $\hat{\mathbf{u}}$ must be computed and there is some extra cost with respect to CORK even in the case the structure of $\mathbf{A} - \lambda \mathbf{B}$ is taken into account for solving the linear system $(\mathbf{A} - \theta_j \mathbf{B}) \hat{\mathbf{u}} = \mathbf{B} \mathbf{u}_j$. On the other hand, there is some overhead cost involved in step 2 of Algorithm 2, since, in CORK, the actual vector \mathbf{u}_j has to be constructed before solving the linear system associated to the shift-and-invert step. Fortunately, according to (2.14), this computation can be arranged as the single matrix-matrix product $Q_j [r_j^{(1)} \ \dots \ r_j^{(d)}]$, where $r_j^{(1)}, \dots, r_j^{(d)}$ are the blocks of the last column of \mathbf{R}_j , which allows optimal efficiency and cache usage on modern computers (see [13, p. 577]).

Inspired in CORK, we will develop in Section 3 the new algorithm R-CORK to solve large-scale and sparse rational eigenvalue problems by using a decomposition similar to (2.14) for the bases of the rational Krylov subspaces associated to the linearization (2.6) of the REP and by working in the spirit of the two levels of orthogonalization originally introduced in TOAR [27, 14]. We will see

that R-CORK has memory and computational advantages similar to those discussed for CORK in the previous paragraph.

3 A new method for solving large-scale and sparse rational eigenvalue problems

3.1 A compact decomposition for rational Krylov subspaces of $\mathcal{A} - \lambda\mathcal{B}$

Consider the matrices \mathcal{A} and \mathcal{B} in (2.7) and the rational Krylov recurrence relation (2.10) which is valid for arbitrary pencils. Our goal is to particularize such relation to the matrices \mathcal{A} and \mathcal{B} in (2.7) in order to save memory and orthogonalization costs. For this purpose, we will partitionate \mathbf{U}_{j+1} conformably to \mathcal{A} and \mathcal{B} as follows

$$\mathbf{U}_{j+1} = [\mathbf{U}_j \quad \mathbf{u}_{j+1}] = \begin{bmatrix} U_j^{(1)} & u_{j+1}^{(1)} \\ U_j^{(2)} & u_{j+1}^{(2)} \\ \vdots & \vdots \\ U_j^{(d)} & u_{j+1}^{(d)} \\ V_j & v_{j+1} \end{bmatrix} \quad (3.1)$$

where $U_j^{(i)} \in \mathbb{C}^{n \times j}$, $u_{j+1}^{(i)} \in \mathbb{C}^n$, for $i = 1, \dots, d$, $V_j \in \mathbb{C}^{s \times j}$, and $v_{j+1} \in \mathbb{C}^s$. Next, following CORK for the first d blocks, we define the matrix $Q_j \in \mathbb{C}^{n \times r_j}$ such that the columns of Q_j are orthonormal with

$$\text{span}\{Q_j\} = \text{span}\{U_j^{(1)}, U_j^{(2)}, \dots, U_j^{(d)}\} \quad (3.2)$$

and $\text{rank}(Q_j) = r_j$. Using (3.2) we can express

$$U_j^{(i)} = Q_j R_j^{(i)}, \quad i = 1, 2, \dots, d, \quad (3.3)$$

where $R_j^{(i)} \in \mathbb{C}^{r_j \times j}$ for $i = 1, 2, \dots, d$. Then, by using (3.3), we have

$$\mathbf{U}_j = \begin{bmatrix} Q_j R_j^{(1)} \\ Q_j R_j^{(2)} \\ \vdots \\ Q_j R_j^{(d)} \\ V_j \end{bmatrix} = \begin{bmatrix} Q_j & & & & \\ & Q_j & & & \\ & & \ddots & & \\ & & & Q_j & \\ & & & & I_s \end{bmatrix} \begin{bmatrix} R_j^{(1)} \\ R_j^{(2)} \\ \vdots \\ R_j^{(d)} \\ V_j \end{bmatrix}. \quad (3.4)$$

By introducing the notation

$$\mathbf{Q}_j := \left[\begin{array}{c|c} (I_d \otimes Q_j) & 0_{dn \times s} \\ \hline 0_{s \times dr_j} & I_s \end{array} \right] \in \mathbb{C}^{(dn+s) \times (dr_j+s)} \quad \text{and} \quad \mathbf{R}_j := \begin{bmatrix} R_j^{(1)} \\ R_j^{(2)} \\ \vdots \\ R_j^{(d)} \\ V_j \end{bmatrix} \in \mathbb{C}^{(dr_j+s) \times j}, \quad (3.5)$$

we have $\mathbf{U}_j = \mathbf{Q}_j \mathbf{R}_j$. Note that the columns of \mathbf{U}_j and \mathbf{Q}_j are orthonormal, so the matrix \mathbf{R}_j has orthonormal columns too. With this notation, we can rewrite (2.10) as the following compact rational Krylov recurrence relation

$$\mathcal{A} \mathbf{Q}_{j+1} \mathbf{R}_{j+1} \underline{H}_j = \mathcal{B} \mathbf{Q}_{j+1} \mathbf{R}_{j+1} \underline{K}_j. \quad (3.6)$$

In order to prove that, as in CORK, we need only one vector to expand Q_j into Q_{j+1} and that, as a consequence, r_j is considerably smaller than jd , i.e., that $\mathbf{Q}_j \mathbf{R}_j$ is indeed a compact representation of \mathbf{U}_j , we will prove first the following Lemmas 3.1 and 3.3. We emphasize the relationship between Lemma 3.1 and Theorem 2.3, but also the difference coming from the presence of the strictly proper part $E(C - \lambda D)^{-1} F^T$ of the rational matrix $R(\lambda)$, which motivates the definition of the rational matrix $\mathbf{A}(\lambda)$ in Lemma 3.1. Apart from this difference, we have stated Lemma 3.1 in a completely analogous way to Theorem 2.3, with the purpose of stressing the relation with CORK, but note that the simple particular structures of M , N , and \mathbf{B} inherited from (2.7)-(2.9) imply that in Lemma 3.1

$$m_0 = 0_{(d-1) \times 1} \quad \text{and} \quad n_0 = -e_{d-1} = - \begin{bmatrix} 0 \\ \vdots \\ 0 \\ 1 \end{bmatrix} \in \mathbb{C}^{d-1},$$

that M_1 and N_1 are also very simple, and that $\bar{\mathbf{B}}_1$ has only one nonzero block. Therefore, the factors $\mathcal{L}(\mu)$ and $\mathcal{U}(\mu)$ in Lemma 3.1 are simpler than the general ones in Theorem 2.3. Note also that Lemma 3.3 is related to [28, Lemma 4.3], although again the strictly proper part of the rational matrix introduces relevant differences.

Lemma 3.1. *Consider a rational matrix*

$$R(\lambda) = P(\lambda) - E(C - \lambda D)^{-1} F^T \in \mathbb{C}(\lambda)^{n \times n},$$

where $P(\lambda) = \sum_{i=0}^d \lambda^i P_i$, $P_i \in \mathbb{C}^{n \times n}$ for $i = 0, \dots, d$, $E, F \in \mathbb{C}^{n \times s}$, $C, D \in \mathbb{C}^{s \times s}$, D is nonsingular, and $E(C - \lambda D)^{-1} F^T$ is a minimal realization. Define the rational matrix

$$\mathbf{A}(\lambda) = \left[\frac{P_{d-1} \quad P_{d-2} \quad \cdots \quad P_1 \quad P_0 - E(C - \lambda D)^{-1} F^T}{M \otimes I_n} \right],$$

and the constant matrix

$$\mathbf{B} = \left[\frac{-P_d \quad 0_n \quad \cdots \quad 0_n \quad 0_n}{N \otimes I_n} \right],$$

with $M = -I_{(d-1) \times d}$ and $N = -[0_{(d-1) \times 1} \quad I_{d-1}]$. Let $\mathcal{P} \in \mathbb{C}^{d \times d}$ be any matrix permutation such that \mathcal{P} moves the first column of a matrix to the last column, this is,

$$\mathcal{P} = [\mathcal{P}_1 \quad e_1], \quad \mathcal{P}_1 \in \mathbb{C}^{d \times (d-1)}, \quad \mathcal{P}_1^T \mathcal{P}_1 = I_{d-1}. \quad (3.7)$$

Then, for every $\mu \in \mathbb{C}$ which is not a pole of $R(\mu)$, i.e., such that $(C - \mu D)$ is nonsingular, we can factorize $\mathbf{A}(\mu) - \mu \mathbf{B}$ as follows

$$\mathbf{A}(\mu) - \mu \mathbf{B} = \mathcal{U}(\mu) \mathcal{L}(\mu) (\mathcal{P} \otimes I_n), \quad (3.8)$$

where

$$\begin{aligned} \mathcal{L}(\mu) &= \begin{bmatrix} R(\mu) & 0 \\ (m_0 - \mu n_0) \otimes I_n & (M_1 - \mu N_1) \otimes I_n \end{bmatrix}, \\ \mathcal{U}(\mu) &= \begin{bmatrix} I_n & (\bar{\mathbf{A}}_1 - \mu \bar{\mathbf{B}}_1) ((M_1 - \mu N_1)^{-1} \otimes I_n) \\ 0 & I_{(d-1)n} \end{bmatrix}, \end{aligned}$$

with

$$M =: [m_0 \quad M_1] \mathcal{P}, \quad N =: [n_0 \quad N_1] \mathcal{P},$$

and

$$\begin{aligned} [P_{d-1} \quad \cdots \quad P_1 \quad P_0 - E(C - \mu D)^{-1} F^T] &=: [P_0 - E(C - \mu D)^{-1} F^T \quad \bar{\mathbf{A}}_1] (\mathcal{P} \otimes I_n), \\ [-P_d \quad 0_n \quad \cdots \quad 0_n] &=: [0_n \quad \bar{\mathbf{B}}_1] (\mathcal{P} \otimes I_n). \end{aligned}$$

Proof. Observe first that the definitions of M , N and \mathcal{P} imply trivially that $M_1 - \mu N_1$ is nonsingular for every $\mu \in \mathbb{C}$. By a direct matrix multiplication, we obtain

$$\begin{aligned} \mathcal{U}(\mu)\mathcal{L}(\mu) &= \begin{bmatrix} R(\mu) + (\bar{\mathbf{A}}_1 - \mu\bar{\mathbf{B}}_1)((M_1 - \mu N_1)^{-1}(m_0 - \mu n_0) \otimes I_n) & (\bar{\mathbf{A}}_1 - \mu\bar{\mathbf{B}}_1) \\ (m_0 - \mu n_0) \otimes I_n & (M_1 - \mu N_1) \otimes I_n \end{bmatrix} \\ &= \left[\frac{R(\mu) + (\bar{\mathbf{A}}_1 - \mu\bar{\mathbf{B}}_1)((M_1 - \mu N_1)^{-1}(m_0 - \mu n_0) \otimes I_n)}{(M - \mu N)\mathcal{P}^T \otimes I_n} \mid (\bar{\mathbf{A}}_1 - \mu\bar{\mathbf{B}}_1) \right]. \end{aligned}$$

Therefore, we only need to prove that

$$R(\mu) + (\bar{\mathbf{A}}_1 - \mu\bar{\mathbf{B}}_1)((M_1 - \mu N_1)^{-1}(m_0 - \mu n_0) \otimes I_n) = P_0 - E(C - \mu D)^{-1}F^T,$$

which is equivalent to prove that

$$P(\mu) + (\bar{\mathbf{A}}_1 - \mu\bar{\mathbf{B}}_1)((M_1 - \mu N_1)^{-1}(m_0 - \mu n_0) \otimes I_n) = P_0. \quad (3.9)$$

The proof of (3.9) is a very simple algebraic manipulation as a consequence of the extremely simple structures of m_0 and n_0 , M_1 and N_1 in this case. Another proof comes from the observation that (3.9) holds because it is proved for proving the ULP decomposition in Theorem 2.3 (see [28, pp. 823-824]). \square

Remark 3.2. Observe that Theorem 2.3 involves the constant $\alpha = e_1^T \mathcal{P}f(\mu)$, which is not present in Lemma 3.1. The reason is that in Lemma 3.1, this constant is equal to 1 as a consequence of the structure of \mathcal{P} and (2.8).

Lemma 3.3. *Let \mathcal{A} and \mathcal{B} be the matrices defined in (2.7). Consider the linear system*

$$(\mathcal{A} - \mu\mathcal{B})\mathbf{x} = \mathcal{B}\mathbf{w}, \quad (3.10)$$

where $\mathbf{x} = [x^{(1)T}, x^{(2)T}, \dots, x^{(d)T}, y^T]^T$ and $\mathbf{w} = [w^{(1)T}, w^{(2)T}, \dots, w^{(d)T}, z^T]^T$, the blocks $x^{(i)}, w^{(i)} \in \mathbb{C}^n$, $i = 1, 2, \dots, d$, $y, z \in \mathbb{C}^s$, and μ is not a pole of $R(\mu)$, i.e., $(C - \mu D)$ is nonsingular. Then, the block $x^{(d)}$ of \mathbf{x} can be computed by solving the following $n \times n$ linear system whose coefficient matrix is $R(\mu)$ in (2.5):

$$R(\mu)x^{(d)} = -P_d w^{(1)} - E(C - \mu D)^{-1}Dz + (\bar{\mathbf{A}}_1 - \mu\bar{\mathbf{B}}_1)((M_1 - \mu N_1)^{-1} \otimes I_n)w^{(2, \dots, d)},$$

where the matrices introduced in Lemma 3.1 are used and $w^{(2, \dots, d)} = [w^{(2)T}, \dots, w^{(d)T}]^T$. The remaining blocks $x^{(i)}$ for $i = 1, \dots, d-1$ of \mathbf{x} can be obtained as linear combinations of $x^{(d)}$ and $w^{(i)}$, $i = 2, \dots, d$. More precisely, if the permutation matrix \mathcal{P} in (3.7) is expressed as $\mathcal{P} = \begin{bmatrix} 0 & 1 \\ \tilde{\mathcal{P}} & 0 \end{bmatrix}$,

then $x^{(1, \dots, d-1)} = [x^{(1)T}, \dots, x^{(d-1)T}]^T$ satisfies the linear system

$$(\tilde{\mathcal{P}} \otimes I_n)x^{(1, \dots, d-1)} = -((M_1 - \mu N_1)^{-1} \otimes I_n)(w^{(2, \dots, d)} + ((m_0 - \mu n_0) \otimes I_n)x^{(d)}).$$

In addition, y can be computed by solving the $s \times s$ linear system

$$(C - \mu D)y = Dz - F^T x^{(d)}.$$

Proof. Rewrite the matrix pencil (2.6) as in (2.9)

$$\mathcal{A} - \mu\mathcal{B} = \left[\begin{array}{c|c} \mathbf{A} - \mu\mathbf{B} & e_1 \otimes E \\ \hline e_d^T \otimes F^T & C - \mu D \end{array} \right]$$

with

$$\mathbf{A} = \left[\begin{array}{cccc} P_{d-1} & P_{d-2} & \cdots & P_0 \\ \hline M \otimes I_n & & & \end{array} \right], \quad \mathbf{B} = \left[\begin{array}{cccc} -P_d & 0_n & \cdots & 0_n \\ \hline N \otimes I_n & & & \end{array} \right],$$

$$M = -I_{(d-1) \times d}, \quad \text{and} \quad N = -[0_{(d-1) \times 1} \quad I_{(d-1)}].$$

Then, we can solve the system (3.10) by solving

$$(\mathbf{A} - \mu\mathbf{B})x^{(1,2,\dots,d)} + (e_1 \otimes E)y = \mathbf{B}w^{(1,2,\dots,d)}, \quad (3.11)$$

$$(e_d^T \otimes F^T)x^{(1,2,\dots,d)} + (C - \mu D)y = Dz, \quad (3.12)$$

where $x^{(1,2,\dots,d)} = [x^{(1)T}, x^{(2)T}, \dots, x^{(d)T}]^T$ and $w^{(1,2,\dots,d)} = [w^{(1)T}, w^{(2)T}, \dots, w^{(d)T}]^T$. The second equation is the equation for y in the statement. By replacing $y = (C - \mu D)^{-1}(Dz - F^T x^{(d)})$ from (3.12) in (3.11), and using the notation of Lemma 3.1, we obtain

$$\begin{aligned} (\mathbf{A} - \mu\mathbf{B})x^{(1,2,\dots,d)} - (e_1 \otimes E)(C - \mu D)^{-1}F^T x^{(d)} &= \mathbf{B}w^{(1,2,\dots,d)} - (e_1 \otimes E)(C - \mu D)^{-1}Dz, \\ \left[\frac{P_{d-1} + \mu P_d \quad P_{d-2} \quad \cdots \quad P_0 - E(C - \mu D)^{-1}F^T}{(M - \mu N) \otimes I_n} \right] x^{(1,2,\dots,d)} &= - \left[\frac{P_d w^{(1)} + E(C - \mu D)^{-1}Dz}{w^{(2,\dots,d)}} \right], \\ (\mathbf{A}(\mu) - \mu\mathbf{B})x^{(1,2,\dots,d)} &= - \left[\frac{P_d w^{(1)} + E(C - \mu D)^{-1}Dz}{w^{(2,\dots,d)}} \right]. \end{aligned}$$

By combining the factorization (3.8) in Lemma 3.1 and the equation above, it is immediate to see that the blocks $x^{(i)}$ for $i = 1, \dots, d-1$ of \mathbf{x} are linear combinations of $x^{(d)}$ and the blocks $w^{(i)}$, $i = 2, \dots, d$. In addition, some elementary matrix manipulations with the matrices $\mathcal{U}(\mu)$ and $\mathcal{L}(\mu)$ in (3.8) and the structure of the permutation matrix \mathcal{P} lead to the equations for $x^{(1,\dots,d-1)}$ and $x^{(d)}$ in the statement. This finishes the proof. \square

Remark 3.4. Since the matrices P_i in the matrix \mathcal{A} in (2.7) appear in decreasing index order, it is natural to choose in Lemmas 3.1 and 3.3 the permutation \mathcal{P} as follows

$$\mathcal{P} = \begin{bmatrix} & & & 1 \\ & & \cdot & \\ & & & \\ 1 & & & \end{bmatrix}. \quad (3.13)$$

As announced before, Lemma 3.3 is the key result that allows us to prove through Theorems 3.6 and 3.7 that only one vector is needed to expand Q_j into Q_{j+1} and, so, that the representation (3.4) for \mathbf{U}_j is indeed compact. Moreover, the equations for $x^{(d)}$, $x^{(1,\dots,d-1)}$, and y deduced in Lemma 3.3 lead to the efficient Algorithm 3 for solving the linear system (3.10), which is fundamental for performing efficiently the shift-and-invert step in the R-CORK method developed in Section 3.2. Observe that in Algorithm 3 a notation similar to that in Lemma 3.3 is used and that $t^{(1,\dots,d)} := (\mathcal{P} \otimes I_n)x^{(1,\dots,d)}$ is computed first, performing later the inverse permutation for getting $x^{(1,\dots,d)}$.

Algorithm 3 Solver for the linear system $(\mathcal{A} - \mu\mathcal{B})\mathbf{x} = \mathcal{B}\mathbf{w}$, with \mathcal{A} and \mathcal{B} as in (2.7)

Input: $\mathcal{A}, \mathcal{B} \in \mathbb{C}^{(nd+s) \times (nd+s)}$ as in (2.7), $\mu \in \mathbb{C}$ such that $(C - \mu D)^{-1}$ exists and $\mathbf{w} \in \mathbb{C}^{nd+s}$.

Output: The solution \mathbf{x} of the linear system.

1. Compute $t = \mathbf{B}w^{(1,2,\dots,d)} - (e_1 \otimes E)(C - \mu D)^{-1}Dz$ as $t = - \left[\frac{P_d w^{(1)} + E(C - \mu D)^{-1}Dz}{w^{(2,\dots,d)}} \right]$.

Solve the block upper triangular system associated to $\mathcal{U}(\mu)$ in (3.8):

2. $t^{(1)} = t^{(1)} - (\bar{\mathbf{A}}_1 - \mu\bar{\mathbf{B}}_1)((M_1 - \mu N_1)^{-1} \otimes I_n)t^{(2,\dots,d)}$.

Solve the block lower triangular system associated to $\mathcal{L}(\mu)$ in (3.8):

3. $t^{(1)} = (R(\mu))^{-1}t^{(1)}$.

4. $t^{(2,\dots,d)} = ((M_1 - \mu N_1)^{-1} \otimes I_n)(t^{(2,\dots,d)} - ((m_0 - \mu n_0) \otimes I_n)t^{(1)})$.

Permute the blocks of $t^{(1,2,\dots,d)}$:

5. $x^{(1,2,\dots,d)} = (\mathcal{P}^T \otimes I_n)t^{(1,2,\dots,d)}$.

Compute the block y of \mathbf{x}

6. $y = (C - \mu D)^{-1}(Dz - F^T x^{(d)})$.

Remark 3.5. The multiplications by inverses in Algorithm 3 have to be understood, in principle, as solutions of linear systems and the key observation on Algorithm 3 is that all the involved linear systems have sizes smaller than the size $(nd + s) \times (nd + s)$ of $(\mathcal{A} - \mu\mathcal{B})$ as we discuss in this remark. The only linear system which is always large is the one in line 3 involving $R(\mu)$ which has the size $n \times n$ of the original REP. Clearly, solving the system in line 3 requires to construct $R(\mu)$ (analogously to [13, 28]), which in the case $(C - \mu D)$ is large and complicated might be performed more efficiently through (1.2) than through (2.5), though this depends on each particular problem. However, we emphasize once again that the matrix $(C - \mu D) \in \mathbb{C}^{s \times s}$ is in many applications [18, 26] very small, since $s \ll n$, and has in addition a very simple structure, which imply that it is often possible just to compute $(C - \mu D)^{-1}$ and to perform the corresponding matrix multiplications to construct $R(\mu)$ through (2.5). These comments on the size $s \ll n$ also apply to the linear systems involving $(C - \mu D) \in \mathbb{C}^{s \times s}$ in lines 1 and 6 which are very often in practice very small. Finally, the linear systems involving $(M_1 - \mu N_1) \otimes I_n$ have size $(d-1)n \times (d-1)n$ and look very large, but they are block linear systems very easy to solve with cost $2n(d-2)$ flops by using simple two term recurrence relations. For instance, if \mathcal{P} is the permutation in (3.13) then

$$M_1 - \mu N_1 = \begin{bmatrix} & & \mu & -1 \\ & \ddots & \ddots & \\ \mu & \ddots & & \\ -1 & & & \end{bmatrix} \in \mathbb{C}^{(d-1) \times (d-1)},$$

and the solution of $((M_1 - \mu N_1) \otimes I_n)\mathbf{x} = \mathbf{b}$ partitioning the vectors in $(d-1)$ blocks of size $n \times 1$ can be obtained as $x^{(1)} = -b^{(d-1)}$ and $x^{(i)} = \mu x^{(i-1)} - b^{(d-i)}$ for $i = 2, 3, \dots, d-1$.

The following theorems are similar to results obtained in [28, Theorems 4.4 and 4.5].

Theorem 3.6. *Let Q_j be defined as in (3.2). Then,*

$$\text{span}\{Q_{j+1}\} = \text{span}\{Q_j, u_{j+1}^{(d)}\}. \quad (3.14)$$

Proof. The proof is immediate from definition (3.2) and Lemma 3.3 with $\mu = \theta_j$ and $\mathbf{w} = \mathbf{u}_j$ (see proof of [28, Theorem 4.4]). \square

Theorem 3.7. *Let Q_j be defined as in (3.2). Then*

$$r_j < d + j. \quad (3.15)$$

Proof. We will prove this theorem by induction. From the definition of Q_j in (3.2), we have that

$$\text{span}\{Q_1\} = \text{span}\{u_1^{(1)}, u_1^{(2)}, \dots, u_1^{(d)}\},$$

so $r_1 \leq d$. Assuming that the inequality (3.15) is satisfied until $j-1$, then we have by Theorem 3.6 that $r_j \leq r_{j-1} + 1 < d + j$. \square

By considering the inequality (3.15) and the fact that r_j increases at most 1 in each iteration, we will show the possible structures of the expansion of the first d blocks of the matrix \mathbf{R}_j defined in (3.5).

Lemma 3.8. *Let $\mathbf{R}_j \in \mathbb{C}^{(dr_j+s) \times j}$ be defined as in (3.5). Then, the first d blocks of the matrix $\mathbf{R}_{j+1} \in \mathbb{C}^{(dr_{j+1}+s) \times (j+1)}$ can take the following forms:*

- if $r_{j+1} > r_j$

$$R_{j+1}^{(i)} = \begin{bmatrix} R_j^{(i)} & r_{j+1}^{(i)} \\ 0_{1 \times j} & \end{bmatrix}, \quad i = 1, 2, \dots, d,$$

where $r_{j+1}^{(i)} \in \mathbb{C}^{r_{j+1}}$, or

- if $r_{j+1} = r_j$

$$R_{j+1}^{(i)} = \begin{bmatrix} R_j^{(i)} & r_{j+1}^{(i)} \end{bmatrix}, \quad i = 1, 2, \dots, d,$$

with $r_{j+1}^{(i)} \in \mathbb{C}^{r_{j+1}}$.

3.2 The R-CORK method

In this section, we will introduce the method to solve large-scale and sparse rational eigenvalue problems based on the compact representation presented in Section 3.1 of the orthonormal bases of the rational Krylov subspaces of the linearization $\mathcal{A} - \lambda\mathcal{B}$ in (2.6). First, we consider an initial vector $\mathbf{u}_1 \in \mathbb{C}^{n+d+s}$ with $\|\mathbf{u}_1\|_2 = 1$ partitioned as in (3.1) and then, we express that vector in a compact form:

$$\mathbf{u}_1 = \begin{bmatrix} u_1^{(1)} \\ \vdots \\ u_1^{(d)} \\ v_1 \end{bmatrix} = \begin{bmatrix} Q_1 R_1^{(1)} \\ \vdots \\ Q_1 R_1^{(d)} \\ v_1 \end{bmatrix}$$

where $Q_1 \in \mathbb{C}^{n \times r_1}$ has orthonormal columns such that

$$\text{span}\{Q_1\} = \text{span}\{u_1^{(1)}, \dots, u_1^{(d)}\}, \quad r_1 = \text{rank}([u_1^{(1)} \dots u_1^{(d)}]).$$

Observe that $r_1 = 1$ if and only if \mathbf{u}_1 is chosen to have collinear nonzero blocks $u_1^{(1)}, \dots, u_1^{(d)}$. Now, taking into account the definition of \mathbf{R}_j in (3.5), after j steps we want to expand Q_j into Q_{j+1} and \mathbf{R}_j into \mathbf{R}_{j+1} , which results in the so-called two levels of orthogonalization.

First level of orthogonalization. In Theorem 3.6 we have proved that we need to orthogonalize $u_{j+1}^{(d)}$ with respect to Q_j to obtain the last orthonormal column of Q_{j+1} . In addition, it can be easily seen that

$$\text{span}\{Q_{j+1}\} = \text{span}\{Q_j, u_{j+1}^{(d)}\} = \text{span}\{Q_j, \hat{u}^{(d)}\}, \quad (3.16)$$

where $\hat{u}^{(d)}$ is the d -th block of size n of the vector $\hat{\mathbf{u}}$ obtained by applying the shift-and-invert step to \mathbf{u}_j (step 2 in Algorithm 1) when $\hat{\mathbf{u}}$ is partitioned as in (3.1). Therefore, we only need to compute the block $\hat{u}^{(d)}$ of $\hat{\mathbf{u}}$ to compute Q_{j+1} . Thus, we can run Algorithm 3 with $\mathbf{w} = \mathbf{u}_j$ and $\mu = \theta_j$ until step 3, saving the resulting vector $t^{(1)} = \hat{u}^{(d)}$. It is important to observe that the first d blocks of \mathbf{u}_j have to be constructed, since the variables in R-CORK are Q_j and \mathbf{R}_j , and \mathbf{u}_j is not stored. As in CORK, they are computed as the single matrix-matrix product $Q_j [r_j^{(1)} \dots r_j^{(d)}]$, where $r_j^{(1)}, \dots, r_j^{(d)}$ are the first d blocks of the last column \mathbf{r}_j of \mathbf{R}_j , which is a very efficient computation in terms of cache utilisation on modern computers. Once $\hat{u}^{(d)}$ is available, by (3.16) we can decompose

$$\hat{u}^{(d)} = Q_j x_j + \alpha_j q_{j+1}, \quad (3.17)$$

where q_{j+1} is a unit vector orthogonal to Q_j and $x_j = Q_j^* \hat{u}^{(d)}$. Observe also that since $\hat{u}^{(d)}$ has been already computed, we can compute the last s entries of $\hat{\mathbf{u}}$, denoted by \hat{v} , from step 6 in Algorithm 3 without the need of performing steps 4 and 5. The vector \hat{v} will be used in the second level of orthogonalization. Now, if $\hat{u}^{(d)}$ does not lie in the subspace spanned by the columns of Q_j , i.e., if $\hat{u}^{(d)} - Q_j x_j \neq 0$, we can expand Q_j into Q_{j+1} as follows:

$$Q_{j+1} = [Q_j \quad q_{j+1}], \quad r_{j+1} = r_j + 1.$$

On the other hand, if $\hat{u}^{(d)}$ lies in the subspace spanned by the columns of Q_j , we have $Q_{j+1} = Q_j$ and $r_{j+1} = r_j$. We summarize the first level of orthogonalization in Algorithm 4. In step 2, if it is necessary, we can reorthogonalize \tilde{q} to ensure orthogonality. In fact, in our MATLAB code, we perform the classical Gram-Schmidt method twice.

Algorithm 4 First level of orthogonalization in R-CORK

Input: The matrix $Q_j \in \mathbb{C}^{n \times r_j}$ and the vector $\hat{u}^{(d)} \in \mathbb{C}^n$ (the d -th block of $\hat{\mathbf{u}} = (\mathcal{A} - \theta_j \mathcal{B})^{-1} \mathcal{B} \mathbf{u}_j$).

Output: The matrix $Q_{j+1} \in \mathbb{C}^{n \times r_{j+1}}$, the vector x_j , and the scalar α_j .

Expanding Q_j into Q_{j+1} .

1. $x_j = Q_j^* \hat{u}^{(d)}$.

2. $\tilde{q} = \hat{u}^{(d)} - Q_j x_j$.

3. $\alpha_j = \|\tilde{q}\|_2$.

if $\alpha_j \neq 0$ **then**

4a. $Q_{j+1} = [Q_j \quad \tilde{q}/\alpha_j]$.

5a. $r_{j+1} = r_j + 1$.

else

4b. $Q_{j+1} = Q_j$.

5b. $r_{j+1} = r_j$.

end if

Second level of orthogonalization. In Algorithm 1, after choosing the shift and performing the shift-and-invert step, we need to compute the entries of the j -th column of \underline{H}_j in step 3. Let us see how to do it efficiently in R-CORK. By using the compact representation of \mathbf{U}_j in (3.4) - (3.5), we have

$$\begin{aligned} h_j &= \mathbf{U}_j^* \hat{\mathbf{u}}, \\ &= (R_j^{(1)})^* Q_j^* \hat{u}^{(1)} + \cdots + (R_j^{(d)})^* Q_j^* \hat{u}^{(d)} + V_j^* \hat{v}, \end{aligned} \quad (3.18)$$

where $\hat{\mathbf{u}}$ has been partitioned in an analogous way to (3.1). Since $(\mathcal{A} - \theta_j \mathcal{B}) \hat{\mathbf{u}} = \mathcal{B} \mathbf{u}_j$, and \mathcal{A} and \mathcal{B} have the structures in (2.7), we obtain the following relation between the blocks of size n of $\hat{\mathbf{u}}$ and the blocks of size n of \mathbf{u}_j

$$\hat{u}^{(i-1)} = \theta_j \hat{u}^{(i)} + u_j^{(i)}, \quad \text{for } i = d, d-1, \dots, 2. \quad (3.19)$$

Motivated by (3.19), we consider the vectors $x_j \in \mathbb{C}^{r_j}$ obtained in step 1 in Algorithm 4 and $\hat{v} \in \mathbb{C}^s$ obtained in step 6 in Algorithm 3 with $x^{(d)} = \hat{u}^{(d)}$ and $\mu = \theta_j$, and a vector $\hat{\mathbf{p}} \in \mathbb{C}^{dr_j+s}$ partitioned as follows

$$\hat{\mathbf{p}} = \begin{bmatrix} \hat{p}^{(1)} \\ \vdots \\ \hat{p}^{(d)} \\ \hat{v} \end{bmatrix}, \quad \hat{p}^{(i)} \in \mathbb{C}^{r_j}, \quad i = 1, \dots, d, \quad (3.20)$$

with the blocks defined by the recurrence relation

$$\begin{aligned} \hat{p}^{(d)} &= x_j, \\ \hat{p}^{(i-1)} &= \theta_j \hat{p}^{(i)} + r_j^{(i)}, \quad i = d, d-1, \dots, 2, \end{aligned} \quad (3.21)$$

where $r_j^{(i)}$ represents the j -th column of the block $R_j^{(i)}$ in (3.4). If $\alpha_j \neq 0$ in step 3 in Algorithm 4, by using $\hat{\mathbf{p}}$, the decomposition (3.17) and the recurrence relation (3.19), the vectors $\hat{u}^{(i)}$, $i = 1, \dots, d$, corresponding to the partition of $\hat{\mathbf{u}}$ as in (3.1) can be represented as follows

$$\hat{u}^{(i)} = Q_{j+1} \begin{bmatrix} \hat{p}^{(i)} \\ \theta_j^{d-i} \alpha_j \end{bmatrix}, \quad i = 1, \dots, d, \quad (3.22)$$

whereas that if $\alpha_j = 0$, we can represent the blocks $\hat{u}^{(i)}$, $i = 1, \dots, d$, as follows

$$\hat{u}^{(i)} = Q_j \hat{p}^{(i)}. \quad (3.23)$$

Then, by using either (3.22) or (3.23) (depending on the value of α_j) in (3.18) and recalling that the columns of Q_{j+1} are orthonormal, we have

$$h_j = \begin{bmatrix} R_j^{(1)} \\ \vdots \\ R_j^{(d)} \\ V_j \end{bmatrix}^* \begin{bmatrix} \hat{p}^{(1)} \\ \vdots \\ \hat{p}^{(d)} \\ \hat{v} \end{bmatrix} = \mathbf{R}_j^* \hat{\mathbf{p}}. \quad (3.24)$$

Thus, after computing $\hat{\mathbf{p}}$ with the recurrence relation (3.21), we can compute h_j by performing a matrix-vector multiplication of size $dr_j + s$, which, according to (3.15), is much smaller than $dn + s$ in large-scale problems and, even more, much smaller than n whenever $s \ll n$ as often happens in applications [18, 26].

Next, in step 4 of Algorithm 1, we need to compute the vector $\tilde{\mathbf{u}}$, which means that in R-CORK we need its compact representation. By using the compact representation of \mathbf{U}_j and (3.22), we have, if $\alpha_j \neq 0$,

$$\begin{aligned} \tilde{\mathbf{u}} &= \hat{\mathbf{u}} - \mathbf{U}_j h_j, \\ &= \begin{bmatrix} Q_{j+1} \begin{bmatrix} \hat{p}^{(1)} \\ \theta_j^{d-1} \alpha_j \end{bmatrix} \\ \vdots \\ Q_{j+1} \begin{bmatrix} \hat{p}^{(d)} \\ \alpha_j \end{bmatrix} \\ \hat{v} \end{bmatrix} - \begin{bmatrix} Q_j R_j^{(1)} \\ \vdots \\ Q_j R_j^{(d)} \\ V_j \end{bmatrix} h_j, \\ &= \begin{bmatrix} Q_{j+1} & & & \\ & \ddots & & \\ & & Q_{j+1} & \\ & & & I_s \end{bmatrix} \begin{bmatrix} \begin{bmatrix} \hat{p}^{(1)} - R_j^{(1)} h_j \\ \theta_j^{d-1} \alpha_j \end{bmatrix} \\ \vdots \\ \begin{bmatrix} \hat{p}^{(d)} - R_j^{(d)} h_j \\ \alpha_j \end{bmatrix} \\ \hat{v} - V_j h_j \end{bmatrix}, \end{aligned}$$

and, in a similar way, if $\alpha_j = 0$ we obtain

$$\tilde{\mathbf{u}} = \begin{bmatrix} Q_j & & & \\ & \ddots & & \\ & & Q_j & \\ & & & I_s \end{bmatrix} \begin{bmatrix} \hat{p}^{(1)} - R_j^{(1)} h_j \\ \vdots \\ \hat{p}^{(d)} - R_j^{(d)} h_j \\ \hat{v} - V_j h_j \end{bmatrix}.$$

Defining

$$\tilde{\mathbf{p}} := \begin{bmatrix} \tilde{p}^{(1)} \\ \vdots \\ \tilde{p}^{(d)} \\ \tilde{v} \end{bmatrix}, \quad \tilde{p}^{(i)} := \hat{p}^{(i)} - R_j^{(i)} h_j \in \mathbb{C}^{r_j}, \quad i = 1, \dots, d, \quad \tilde{v} := \hat{v} - V_j h_j \in \mathbb{C}^s, \quad (3.25)$$

and taking into account that the columns of Q_{j+1} and Q_j are orthonormal, we can express the step

5 in Algorithm 1 as follows: if $\alpha_j \neq 0$, then

$$h_{j+1,j} = \|\tilde{\mathbf{u}}\|_2 = \left\| \begin{bmatrix} \tilde{p}^{(1)} \\ \theta_j^{d-1} \alpha_j \\ \vdots \\ \tilde{p}^{(d)} \\ \alpha_j \\ \tilde{v} \end{bmatrix} \right\|_2 \quad \text{and} \quad \mathbf{u}_{j+1} = \mathbf{Q}_{j+1} \cdot \frac{1}{h_{j+1,j}} \begin{bmatrix} \tilde{p}^{(1)} \\ \theta_j^{d-1} \alpha_j \\ \vdots \\ \tilde{p}^{(d)} \\ \alpha_j \\ \tilde{v} \end{bmatrix}, \quad (3.26)$$

where the notation in (3.5) is used, while if $\alpha_j = 0$ we proceed as in (3.26) by removing all the entries involving α_j and with $\mathbf{Q}_{j+1} = \mathbf{Q}_j$. From the previous equations, we can conclude that the first d blocks of size r_{j+1} of the last column of \mathbf{R}_{j+1} in (3.5) are given if $\alpha_j \neq 0$ by

$$r_{j+1}^{(i)} = \frac{1}{h_{j+1,j}} \begin{bmatrix} \tilde{p}^{(i)} \\ \theta_j^{d-i} \alpha_j \end{bmatrix}, \quad i = 1, \dots, d, \quad (3.27)$$

and if $\alpha_j = 0$ by

$$r_{j+1}^{(i)} = \frac{1}{h_{j+1,j}} \tilde{p}^{(i)}, \quad i = 1, \dots, d. \quad (3.28)$$

In addition, the last block of size s of the last column of \mathbf{R}_{j+1} is

$$v_{j+1} = \frac{1}{h_{j+1,j}} \tilde{v}. \quad (3.29)$$

Since \mathbf{R}_j has orthonormal columns, from (3.24) and the definitions in (3.20)-(3.21) and (3.25), we have that h_j and $\tilde{\mathbf{p}}$ satisfy

$$\tilde{\mathbf{p}} = \hat{\mathbf{p}} - \mathbf{R}_j h_j, \quad (3.30)$$

where $\tilde{\mathbf{p}}$ is orthogonal to \mathbf{R}_j . This process is the Gram-Schmidt process without the normalization step, and it is summarized in Algorithm 5.

Algorithm 5 Second level of orthogonalization in R-CORK

Input: The matrix $\mathbf{R}_j \in \mathbb{C}^{(dr_j+s) \times j}$ and the vector $\hat{\mathbf{p}} \in \mathbb{C}^{dr_j+s}$ from (3.20)-(3.21).

Output: Vectors $h_j \in \mathbb{C}^j$ and $\tilde{\mathbf{p}} \in \mathbb{C}^{dr_j+s}$.

1. $h_j = \mathbf{R}_j^* \hat{\mathbf{p}}$.
 2. $\tilde{\mathbf{p}} = \hat{\mathbf{p}} - \mathbf{R}_j h_j$.
-

Remark 3.9. In order to improve orthogonality, a reorthogonalization method can be included in Algorithm 5. In our MATLAB code, we use the classical Gram-Schmidt process twice.

The whole procedure of this new method to solve large-scale and sparse rational eigenvalue problems requires the use of the two levels of orthogonalization described in this section, the first level to expand Q_j into Q_{j+1} and the second level to expand \mathbf{R}_j into \mathbf{R}_{j+1} . The complete R-CORK method is summarized in Algorithm 6. Note that R-CORK has as inputs the matrix Q_1 and the vector \mathbf{R}_1 , which have to be computed. As in CORK [28, p. 830], there are two possible ways of computing these inputs: either starting with a random vector $\mathbf{u}_1 \in \mathbb{C}^{nd+s}$ and using an economy-size QR factorization, or emulating the structure of the eigenvectors (2.8) of the linearization in (2.6). The details are very similar to the ones in [28, p. 830] and are omitted. Recall in Algorithm 6 that \mathbf{r}_j denotes the last column of the matrix \mathbf{R}_j in (3.5).

Algorithm 6 Compact rational Krylov method for REP (R-CORK)

Input: $Q_1 \in \mathbb{C}^{n \times r_1}$, $\mathbf{R}_1 \in \mathbb{C}^{(dr_1+s) \times 1}$ with $Q_1^* Q_1 = I_{r_1}$ and $\mathbf{R}_1^* \mathbf{R}_1 = 1$.

Output: Approximate eigenpairs (λ, \mathbf{x}) of $\mathcal{A} - \lambda \mathcal{B}$ with \mathcal{A} and \mathcal{B} as in (2.7).

for $j = 1, 2, \dots$ **do**

1. Choose shift θ_j .

2. Compute $\mathbf{u}_j = \mathbf{Q}_j \mathbf{r}_j$, obtaining the first d blocks as matrix-matrix product $Q_j [r_j^{(1)} \dots r_j^{(d)}]$.

3. Compute $\hat{u}^{(d)}$ by using Algorithm 3 until step 3 applied to $\mathbf{w} = \mathbf{u}_j$ and $\mu = \theta_j$.

4. Compute \hat{v} from step 6 in Algorithm 3.

First level of orthogonalization

5. Run Algorithm 4 obtaining Q_{j+1} , the scalar α_j and the vector x_j .

Second level of orthogonalization:

6. Compute $\hat{\mathbf{p}}$ in (3.20) via the recurrence relation in (3.21).

7. Run Algorithm 5 obtaining $\tilde{\mathbf{p}}$ and h_j .

8. Compute $h_{j+1,j}$ and \mathbf{r}_{j+1} using (3.26)-(3.27)-(3.28)-(3.29) and get \mathbf{R}_{j+1} with Lemma 3.8.

9. Compute eigenpairs: (λ_i, t_i) of (2.12) and test for convergence.

end for

10. Compute eigenvectors: $\mathbf{x}_i = \mathbf{Q}_{j+1} \mathbf{R}_{j+1} \underline{H}_j t_i$.

3.3 Memory and computational costs

In this section, we discuss the memory and the computational costs of R-CORK and compare these costs with those of the classical rational Krylov (RK) method, i.e., Algorithm 1, applied directly to the linearization $\mathcal{A} - \lambda \mathcal{B}$ of the REP in (2.6). In order to simplify the results we will take $r_j = j + d$ in R-CORK, which is the upper bound in Theorem 3.7 and that essentially corresponds to start the R-CORK iteration with $Q_1 \in \mathbb{C}^{n \times d}$ ($r_1 = d$) or, equivalently, with a random initial vector \mathbf{u}_1 whose first d blocks in the partition (3.1) are linearly independent. If the first d blocks of \mathbf{u}_1 are taken to be collinear, then one can take $r_j = j$ and to improve even more the costs of R-CORK. In addition, note that we estimate the costs for any value of s , where $s \times s$ is the size of the lower-right block $C - \lambda D$ of $\mathcal{A} - \lambda \mathcal{B}$ appearing in the strictly proper part of the REP (2.5). In this way, it will be seen that even if $s \approx n$, R-CORK has considerable advantages with respect to RK in terms of memory and computational costs. However, we emphasize that such advantages are still much more relevant when $s \ll n$, as happens very often in applications [18, 26].

For the memory costs, after j iterations R-CORK stores $Q_j \in \mathbb{C}^{n \times r_j}$ and $\mathbf{R}_j \in \mathbb{C}^{(dr_j+s) \times j}$, which amounts to $(n + dj)(j + d) + sj \approx n(j + d) + sj$ numbers. Note that the approximation $n + dj \approx n$ holds in any reasonable large-scale REP. In contrast, RK stores \mathbf{U}_j , which amounts to $(nd + s)j = ndj + sj$ numbers. Since, $(j + d) < dj$ for most reasonable choices of j and degrees d appearing in practice, we see that R-CORK is much more memory-efficient than RK. These memory costs are shown in Table 3.1.

With respect to the computational costs, observe that for both R-CORK and RK the cost is the sum of (i) the shift-and-invert step and (ii) the orthogonalization steps. Let us analyze first the shift-and-invert steps. If the shift-and-invert step in RK, i.e., step 2 in Algorithm 1, is performed by applying an unstructured solver to the $(nd + s) \times (nd + s)$ linear system $(\mathcal{A} - \theta_j \mathcal{B}) \hat{\mathbf{u}} = \mathcal{B} \mathbf{u}_j$, then the cost of RK is much larger than the cost of R-CORK, since R-CORK solves this system with Algorithm 3 (removing step 4) which is much more efficient because requires the solution of smaller linear systems (essentially, see Remark 3.5, one of size $n \times n$ and two of size $s \times s$, which are very often extremely small). However, one can consider to perform the shift-and-invert step in RK with Algorithm 3, but this is still somewhat more expensive than R-CORK, because for RK it is needed to perform step 4 of Algorithm 3, with an additional cost of $2n(d - 2)$ flops in each iteration (see Remark 3.5). A final important remark on the shift-and-invert step is that R-CORK involves the

overhead cost of constructing \mathbf{u}_j in step 2 of Algorithm 6, which in RK is not needed. However, note that, as explained in previous sections, this construction can be performed as in CORK via a single matrix-matrix product, which allows for optimal efficiency and cache utilisation on modern computers [13, p. 577]. Moreover, we emphasize that a traditional construction of \mathbf{u}_j in R-CORK costs $\mathcal{O}(dnr_j) = \mathcal{O}(dn(j+d)) \approx \mathcal{O}(dnj)$ flops at iteration j , which added to the orthogonalization cost of R-CORK discussed below would give a cost of the same order of the orthogonalization cost of RK.

Finally, we discuss the orthogonalization costs of RK and R-CORK. In RK, the orthogonalization is performed in steps 3-4-5 of Algorithm 1 and its cost is well-known to be $\mathcal{O}(j(nd+s)) = \mathcal{O}(jnd+js)$ flops at iteration j , which amounts to $\mathcal{O}(j^2nd+j^2s)$ flops in the first j iterations (see Table 3.1). In R-CORK, the orthogonalization is performed in steps 5-6-7-8 of Algorithm 6. At iteration j , the cost of step 5 is $\mathcal{O}(r_jn) = \mathcal{O}((j+d)n)$ flops, the cost of step 6 is $\mathcal{O}(r_jd) = \mathcal{O}((j+d)d)$ flops, which is negligible with respect to the cost of step 5, the cost of step 7 is $\mathcal{O}(j(dr_j+s)) = \mathcal{O}(jd(j+d)+js)$ flops, and the cost of step 8 is $\mathcal{O}(dr_j+s) = \mathcal{O}(d(j+d)+s)$ flops. Therefore, the total cost at iteration j of the orthogonalization in R-CORK is $\mathcal{O}((n+jd)(j+d)+js) \approx \mathcal{O}(n(j+d)+js)$, where we have used again the approximation $n+jd \approx n$, which gives $\mathcal{O}(j^2n+jdn+j^2s)$ flops in the first j iterations (see Table 3.1). Observe that the orthogonalization cost of RK includes the large term j^2nd which is not present in the cost of R-CORK. Therefore, the orthogonalization cost of R-CORK is considerably smaller than the one of RK.

In Table 3.1, we summarize the comparison of the costs between R-CORK and RK.

	Classical rational Krylov method	R-CORK method
Orthogonalization cost	$\mathcal{O}(j^2nd+j^2s)$	$\mathcal{O}(j^2n+jdn+j^2s)$
Memory cost	$ndj+sj$	$n(j+d)+sj$

Table 3.1: Orthogonalization and memory costs for classical rational Krylov method and R-CORK method after j iterations.

4 Implicit restarting in R-CORK

Practical implementations of any Krylov-type method for computing eigenvalues of large-scale problems require effective restarting strategies. The goal of this section is to develop an implicit restarting strategy for R-CORK that restarts both Q_j and \mathbf{R}_j in the compact representation of \mathbf{U}_j in (3.4)-(3.5). Since R-CORK shares many of the properties of CORK, the results of this section are similar to those in [28, Section 6], which in turn are based on implicit restarting procedures for classical rational Krylov methods [6] and on the Krylov-Schur restart developed for TOAR in [13, Section 4.2].

Following the Krylov-Schur spirit [25] (see also [24, Section 5.2]), the restarting technique we propose transforms first the matrices \underline{H}_j and \underline{K}_j in (3.6) to (quasi)triangular form, in order to reorder the Ritz values and to preserve the desired ones with a rational Krylov subspace of smaller dimension. Second, by representing the new smaller Krylov subspace in its compact form in an efficient way, the implicit restart of R-CORK is completed. The main difference of the process described below with respect to the implicit restarting in [28, Section 6] is that here we need to add a new block of size $s \times s$ corresponding to the rational part of $R(\lambda)$ in (2.5).

Suppose that after j iterations, we have the rational Krylov recurrence relation in its compact form as in (3.6)

$$\mathcal{A}\mathbf{Q}_{j+1}\mathbf{R}_{j+1}\underline{H}_j = \mathcal{B}\mathbf{Q}_{j+1}\mathbf{R}_{j+1}\underline{K}_j, \quad (4.1)$$

and we want to reduce this representation to a smaller compact rational decomposition of size p ,

$p < j$, this is

$$\mathcal{A}\mathbf{Q}_{p+1}^+ \mathbf{R}_{p+1}^+ \underline{H}_p^+ = \mathcal{B}\mathbf{Q}_{p+1}^+ \mathbf{R}_{p+1}^+ \underline{K}_p^+, \quad p < j.$$

For this purpose, we consider the generalized Schur decomposition:

$$H_j = \begin{bmatrix} Y_p & Y_{j-p} \end{bmatrix} \begin{bmatrix} T_{p \times p}^{(H)} & * \\ 0 & T_{(j-p) \times (j-p)}^{(H)} \end{bmatrix} \begin{bmatrix} Z_p^* \\ Z_{j-p}^* \end{bmatrix}, \quad (4.2)$$

$$K_j = \begin{bmatrix} Y_p & Y_{j-p} \end{bmatrix} \begin{bmatrix} T_{p \times p}^{(K)} & * \\ 0 & T_{(j-p) \times (j-p)}^{(K)} \end{bmatrix} \begin{bmatrix} Z_p^* \\ Z_{j-p}^* \end{bmatrix}, \quad (4.3)$$

where H_j and K_j are the $j \times j$ upper Hessenberg matrices obtained by removing the last row of \underline{H}_j and \underline{K}_j respectively, $Y := \begin{bmatrix} Y_p & Y_{j-p} \end{bmatrix}$, $Z := \begin{bmatrix} Z_p & Z_{j-p} \end{bmatrix} \in \mathbb{C}^{j \times j}$ are unitary matrices with $Y_p, Z_p \in \mathbb{C}^{j \times p}$, $Y_{j-p}, Z_{j-p} \in \mathbb{C}^{j \times (j-p)}$ and $T^{(H)} := \begin{bmatrix} T_{p \times p}^{(H)} & * \\ 0 & T_{(j-p) \times (j-p)}^{(H)} \end{bmatrix}$, $T^{(K)} := \begin{bmatrix} T_{p \times p}^{(K)} & * \\ 0 & T_{(j-p) \times (j-p)}^{(K)} \end{bmatrix} \in \mathbb{C}^{j \times j}$ are upper (quasi)triangular matrices with $T_{p \times p}^{(H)}, T_{p \times p}^{(K)} \in \mathbb{C}^{p \times p}$ and $T_{(j-p) \times (j-p)}^{(H)}, T_{(j-p) \times (j-p)}^{(K)} \in \mathbb{C}^{(j-p) \times (j-p)}$. The $p < j$ Ritz values of interest are the eigenvalues of the pencil $T_{p \times p}^{(K)} - \lambda T_{p \times p}^{(H)}$. By multiplying by Z on the right the recurrence relation (4.1) and using (4.2) and (4.3), and considering the first p columns, we obtain:

$$\mathcal{A}\mathbf{Q}_{j+1} \mathbf{R}_{j+1} \begin{bmatrix} Y_p & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} T_{p \times p}^{(H)} \\ h_{j+1,j} \tilde{z}^* \end{bmatrix} = \mathcal{B}\mathbf{Q}_{j+1} \mathbf{R}_{j+1} \begin{bmatrix} Y_p & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} T_{p \times p}^{(K)} \\ k_{j+1,j} \tilde{z}^* \end{bmatrix}, \quad (4.4)$$

where \tilde{z}^* represents the first p entries of the last row of Z . By introducing the notation:

$$Y_1 := \begin{bmatrix} Y_p & 0 \\ 0 & 1 \end{bmatrix} \in \mathbb{C}^{(j+1) \times (p+1)}, \quad \underline{H}_p^+ := \begin{bmatrix} T_{p \times p}^{(H)} \\ h_{j+1,j} \tilde{z}^* \end{bmatrix}, \quad \underline{K}_p^+ := \begin{bmatrix} T_{p \times p}^{(K)} \\ k_{j+1,j} \tilde{z}^* \end{bmatrix} \in \mathbb{C}^{(p+1) \times p}, \quad (4.5)$$

and defining $\mathbf{W}_{p+1} = \mathbf{R}_{j+1} Y_1$, we obtain

$$\mathcal{A}\mathbf{Q}_{j+1} \mathbf{W}_{p+1} \underline{H}_p^+ = \mathcal{B}\mathbf{Q}_{j+1} \mathbf{W}_{p+1} \underline{K}_p^+. \quad (4.6)$$

Note that with this transformation, we reduce the size of the matrices $\underline{H}_p^+, \underline{K}_p^+$, and \mathbf{W}_{p+1} with respect to $\underline{H}_j, \underline{K}_j$, and \mathbf{R}_{j+1} , and remove the Ritz values that are not of interest. However, observe that the large factor \mathbf{Q}_{j+1} remains unchanged. In order to reduce the size of \mathbf{Q}_{j+1} , consider

$$\mathbf{W}_{p+1} = \begin{bmatrix} W_{p+1}^{(1)} \\ \vdots \\ W_{p+1}^{(d)} \\ V_{j+1} Y_1 \end{bmatrix}, \quad W_{p+1}^{(i)} \in \mathbb{C}^{r_{j+1} \times (p+1)}, \quad i = 1, \dots, d,$$

and let ω be the rank of $[W_{p+1}^{(1)} \ \dots \ W_{p+1}^{(d)}]$. The key observation is that although the matrices $\underline{H}_p^+, \underline{K}_p^+$ are no longer in Hessenberg form, the subspace spanned by the columns of $(\mathbf{Q}_{j+1} \mathbf{W}_{p+1})$ is still a rational Krylov subspace corresponding to $\mathcal{A} - \lambda \mathcal{B}$ [6]. Therefore, we can apply Theorem 3.7 to $\text{span} \left\{ Q_{j+1} W_{p+1}^{(1)}, \dots, Q_{j+1} W_{p+1}^{(d)} \right\} = Q_{j+1} \text{span} \left\{ W_{p+1}^{(1)}, \dots, W_{p+1}^{(d)} \right\}$ to obtain that $\omega \leq d+p < d+j$. Then, we compute the economy singular value decomposition of:

$$[W_{p+1}^{(1)} \ \dots \ W_{p+1}^{(d)}] = \mathcal{U} \mathcal{S} [\mathcal{V}^{(1)} \ \dots \ \mathcal{V}^{(d)}],$$

where $\mathcal{U} \in \mathbb{C}^{r_{j+1} \times \omega}$, $\mathcal{S} \in \mathbb{C}^{\omega \times \omega}$ and $\mathcal{V}^{(i)} \in \mathbb{C}^{\omega \times (p+1)}$ for $i = 1, \dots, d$. Thus, by defining

$$Q_{p+1}^+ = Q_{j+1}\mathcal{U}, \quad \mathbf{R}_{p+1}^+ = \begin{bmatrix} \mathcal{S}\mathcal{V}^{(1)} \\ \vdots \\ \mathcal{S}\mathcal{V}^{(d)} \\ V_{j+1}Y_1 \end{bmatrix}, \quad \mathbf{Q}_{p+1}^+ = \begin{bmatrix} Q_{p+1}^+ & & & \\ & \ddots & & \\ & & Q_{p+1}^+ & \\ & & & I_s \end{bmatrix},$$

we get from (4.6) the compact rational Krylov recurrence relation

$$\mathbf{A}\mathbf{Q}_{p+1}^+ \mathbf{R}_{p+1}^+ \underline{H}_p^+ = \mathbf{B}\mathbf{Q}_{p+1}^+ \mathbf{R}_{p+1}^+ \underline{K}_p^+, \quad (4.7)$$

with $p < j$. It is important to emphasize again that the matrices \underline{H}_p^+ and \underline{K}_p^+ are no longer upper Hessenberg matrices, however, they contain the required Ritz values and the columns of $\mathbf{Q}_{p+1}^+ \mathbf{R}_{p+1}^+$ span a corresponding rational Krylov subspace. We continue the process by expanding (4.7) with Algorithm 6 until we get a rational Krylov subspace of dimension j . The matrices \underline{H}_j^+ and \underline{K}_j^+ obtained in this expansion are not in Hessenberg form, although their columns $p+1, \dots, j$ have a Hessenberg structure (see [24, p. 329]). Then, the restarting process described in this section is applied again to get a new compact relation (4.7) of “size p ”. This expansion-restarting procedure is cyclicly repeated until the prescribed stopping criterion is satisfied for a certain desired number, less than or equal to p , Ritz pairs.

5 Numerical experiments

In this section, we present two large-scale and sparse numerical examples to illustrate the efficiency of the R-CORK method. All reported experiments were performed using Matlab R2013a on a PC with a 2,2 GHz Intel (R) Core (TM) i7 processor, with 16 GB of RAM and DDR3 memory type, and with operating system macOS Sierra, version 10.12.1.

By following [28, Section 8], in the numerical experiments we plot the residuals at each iteration, with and without restarts, obtained by using the R-CORK method, the dimension of the subspace at each iteration for R-CORK, and the comparison of the memory storages of R-CORK and of the classical rational Krylov method applied directly to the linearization (2.6). We also report on the number of iterations until convergence.

Inspired by the applications in [26, Section 4], we construct numerical experiments with prescribed eigenvalues and poles of a rational matrix $R(\lambda)$ represented as in (2.5). In order to measure the convergence of an approximate eigenpair (λ, x) of $R(\lambda)$, we consider the relative norm residual:

$$E(\lambda, x) = \frac{\|R(\lambda)x\|_2}{(\sum_{i=0}^d |\lambda|^i \|P_i\|_F + \|E(C - \lambda D)^{-1} F^T\|_F) \|x\|_2}. \quad (5.1)$$

Observe that the computation of $E(\lambda, x)$ involves matrices and vectors of size n and, so, is expensive. Therefore, in actual practice, we recommend to test first the convergence through a cheap estimation of the residual of the linearized problem, i.e., $\|(\mathcal{A} - \lambda\mathcal{B})\mathbf{z}\|_2$, involving only the small projected problem (2.12), and once such residual is sufficiently small to compute the residual (5.1) every $q > 1$ iterations instead of at each iteration. However, in our examples, we performed the computation of $E(\lambda, x)$ at each iteration for the purpose of illustration.

The computation of (5.1) deserves some comments. Note first that it requires to recover the approximated eigenvector x of $R(\lambda)$ from the approximated eigenvector \mathbf{z} of the linearization $\mathcal{A} - \lambda\mathcal{B}$ in (2.6) computed in step 10 of Algorithm 6. This recovery, according to the first equation in (2.8), can be done by taking any of the first d blocks of \mathbf{z} if $\lambda \neq 0$. Since in our numerical examples the moduli of the approximate eigenvalues are larger than 1, we have chosen the first block of \mathbf{z} as approximate x . However, we recommend to choose the d th block if the moduli of the approximate

eigenvalues are smaller than 1. The calculation of the quantities $\|P_i\|_F$, $i = 0, \dots, d$ needs to be performed only once and it is inexpensive since the matrices P_i are sparse in practice. Finally, to compute the expression $\|E(C - \lambda D)^{-1}F^T\|_F$ on the denominator in (5.1), we use

$$\begin{aligned}\|E(C - \lambda D)^{-1}F^T\|_F^2 &= \text{trace}((E(C - \lambda D)^{-1}F^T)^*E(C - \lambda D)^{-1}F^T), \\ &= \text{trace}((E^*E)(C - \lambda D)^{-1}(F^T\bar{F})(C - \lambda D)^{-*}),\end{aligned}$$

which only involves the matrices E^*E , $F^T\bar{F}$, $(C - \lambda D)^{-1}$ and $(C - \lambda D)^{-*}$ of size $s \times s$. Since in many application $s \ll n$, this computation is usually inexpensive.

Numerical experiment 5.1. We construct a REP of the type arising from the free vibrations of a structure if one uses a viscoelastic constitutive relation to describe the behavior of a material [18, 26]. The REPs of this type have the following structure:

$$R(\lambda)x = \left(\lambda^2 M + K - \sum_{i=1}^k \frac{1}{1 + b_i \lambda} \Delta G_i \right) x = 0, \quad (5.2)$$

where the mass and stiffness matrices M and K are real symmetric and positive definite, b_j are relaxation parameters over k regions, and ΔG_j is an assemblage of element stiffness matrices over the region with the distinct relaxation parameters. As in [26], we consider the case where $\Delta G_i = E_i E_i^T$ and $E_i \in \mathbb{R}^{n \times s_i}$. By defining

$$E = [E_1, E_2, \dots, E_k], \quad D = \text{diag}(b_1 I_{s_1}, b_2 I_{s_2}, \dots, b_k I_{s_k}),$$

the REP (5.2) can be written in the form (2.5):

$$(\lambda^2 M + K - E(I + \lambda D)^{-1}E^T)x = 0.$$

In our particular example, we consider the case with one region and one relaxation parameter $b_1 = -1$. The construction of the matrices M and K in our example proceeds as follows: construct first $R_1(\lambda) = \lambda^2 A_2 + A_0 - e_{10000}(1 - \lambda)^{-1}(e_{10000})^T$, with $A_2, A_0 \in \mathbb{R}^{10000 \times 10000}$ diagonal and positive definite matrices and e_{10000} the last column of I_{10000} . This structure allows to prescribe easily the eigenvalues for $R_1(\lambda)$. Then, we consider the following invertible tridiagonal matrix P

$$P = \begin{bmatrix} 1 & \frac{1}{2} & & & \\ \frac{1}{3} & 1 & \ddots & & \\ & \ddots & \ddots & \ddots & \\ & & & \frac{1}{3} & \frac{1}{2} \\ & & & & 1 \end{bmatrix},$$

and finally construct $R(\lambda) = P R_1(\lambda) P^T$. Since P is invertible, the eigenvalues of $R(\lambda)$ and $R_1(\lambda)$ are the same. By using this procedure, we have constructed the REP

$$R(\lambda)x = (\lambda^2 M + K - p_{10000}(1 - \lambda)^{-1}(p_{10000})^T) x = 0, \quad (5.3)$$

where $M := P A_2 P^T$, $K := P A_0 P^T \in \mathbb{R}^{10000 \times 10000}$ are symmetric, positive definite, and pentadiagonal matrices, and $p_{10000} \in \mathbb{R}^{10000}$ represents the last column of the matrix P .

In this example we are interested in computing the 20 eigenvalues of (5.3) with negative imaginary part and with largest absolute value of the negative imaginary part. To aim our goal, we use 3 cyclically repeated shifts in the rational Krylov steps and a random unit real vector as an initial vector. The reader can see the approximate eigenvalues computed by R-CORK and the chosen shifts in Figure 5.1(a). We first solve the REP (5.3) by using Algorithm 6 without restart, and after 85 iterations, we find the required eigenvalues with a tolerance (5.1) of 10^{-10} . The convergence

history is shown in Figure 5.1(b). In Figure 5.1(d), we plot r_j , the rank of Q_j at the iteration j , and j , the dimension of the Krylov subspace. Since we did not perform restart, we can see that both, r_j and j increases with the iteration count j and that $r_j = j + 1$, as expected since the degree of the polynomial part of (5.3) is $d = 2$. Figure 5.1(f) displays the comparison between the cost of memory storage of both the R-CORK method, by using Algorithm 6, and the classical rational Krylov method, by using Algorithm 1. From this figure, we can see that the R-CORK method requires approximately half of the memory storage that the classical rational Krylov method, which is consistent with the degree 2 of the polynomial part of (5.3).

Next, we apply Algorithm 6 to the REP (5.3) combined with the implicit restarting introduced in Section 4. We choose the maximum dimension of the subspace $m = 45$, which is reduced after each restart to dimension $p = 30$ to compute the 20 required eigenvalues. The convergence history of the eigenpairs (λ, x) computed by this restarted R-CORK method is shown in Figure 5.1(c). After 3 restarts and 81 iterations, the 20 required eigenvalues have been found with a tolerance (5.1) of 10^{-10} . In Figure 5.1(e) the reader can see the rank of Q_j at the j -th iteration and it can be seen that with restart, the relation between j and r_j continues the same. Finally, in Figure 5.1(g) we plot the memory storage for R-CORK and classical rational Krylov, and it can be observed that the memory cost for the R-CORK method is a factor close to 2 smaller than the memory cost obtained by the classical rational Krylov method.

Numerical experiment 5.2. For this numerical example, we consider an academic REP of size 5000×5000 and with the degree of its polynomial part equal to 3, i.e., a REP of the form

$$R(\lambda) = \lambda^3 A_3 + \lambda^2 A_2 + \lambda A_1 + A_0 - E(C - \lambda D)^{-1} F^T. \quad (5.4)$$

The coefficient matrices of $R(\lambda)$ in (5.4) were constructed in a similar way that in the numerical experiment 5.1: first, we consider a rational matrix $R_2(\lambda) = \lambda^3 P_3 + \lambda^2 P_2 + \lambda P_1 + P_0 - E_0(C - \lambda D)^{-1} F_0^T$ with prescribed eigenvalues, where $P_i \in \mathbb{R}^{5000 \times 5000}$ are diagonal matrices, $E_0 = [e_1 + e_2, e_5 + e_6]$, $F_0 = [e_{4997} + e_{4998}, e_{4999} + e_{5000}] \in \mathbb{R}^{5000 \times 2}$, with e_i the i th canonical vector of size 5000, and


$$C = \begin{bmatrix} 105 & 0 \\ 0 & -105 \end{bmatrix}, \quad D = I_2,$$

and then we define $R(\lambda) = P R_2(\lambda) Q$, where


$$P = \begin{bmatrix} 1 & \frac{1}{2} & \frac{1}{3} & & & \\ -\frac{1}{4} & \ddots & \ddots & \ddots & & \\ -\frac{1}{5} & \ddots & \ddots & \ddots & \frac{1}{3} & \\ & \ddots & \ddots & \ddots & \frac{1}{2} & \\ & & -\frac{1}{5} & -\frac{1}{4} & 1 & \end{bmatrix}, \quad Q = \begin{bmatrix} -1 & -\frac{1}{3} & & & & \\ \frac{1}{2} & \ddots & \ddots & & & \\ & \ddots & \ddots & & & \\ & & \ddots & & -\frac{1}{3} & \\ & & & \frac{1}{2} & -1 & \end{bmatrix} \in \mathbb{R}^{5000 \times 5000}.$$

The goal of this example is to compute the 30 eigenvalues closest to zero. In this situation, it is natural to choose zero as a fixed shift. In Figure 5.2(a), the approximate eigenvalues computed by R-CORK are displayed. By starting with a random unit complex vector, first we apply R-CORK without restarting, and after 83 iterations, the desired eigenvalues are obtained with a tolerance (5.1) of 10^{-12} . The convergence history can be seen in Figure 5.2(b). In Figure 5.2(d), we see that the relation $r_j < j + d$ with j the number of iterations also holds in this example, though in this case with $d = 3$ since this is the degree of the polynomial part in (5.4). Figure 5.2(f) shows the memory costs of R-CORK and classical rational Krylov. It is observed that the reduction in cost of R-CORK is approximately a factor of 3, i.e., the degree of the polynomial part of (5.4).


As a final example, we solve (5.4) by using R-CORK combined with restarting and taking a maximum subspace dimension $m = 60$ which is reduced to $p = 40$ after every restart. The




(a) Eigenvalues.




(b) Convergence history without restart.




(c) Convergence history with restart.




(d) Dimension of the subspace without restart.



(e) Dimension of the subspace with restart.



(f) Memory cost without restart.



(g) Memory cost with restart.

Figure 5.1: Numerical experiment 5.1.

convergence history is shown in Figure 5.2(c), where it is observed that after 91 iterations and 2 restarts, the 30 eigenvalues closest to zero have been found with a tolerance (5.1) of 10^{-12} . Despite the fact that a few more iterations are needed with restart than without restart, we see in Figure 5.2(e) that we are using a Krylov subspace of much smaller dimension than without restart to compute the eigenpairs. In addition, we emphasize that Figure 5.2(e) shows that for this particular example $r_j < j$ after the restarts, which illustrates that in practice the upper bound $d + j - 1$ in Theorem 3.7 is not always attained and that the memory efficiency of R-CORK can be larger than the one theoretically expected. Finally, the comparison of the memory costs for the R-CORK and for the classical rational Krylov methods is plotted in Figure 5.2(g), where we see again that the cost of R-CORK is approximately a factor $d = 3$ smaller.


6 Conclusions and lines of future research

In this paper, we have introduced the R-CORK method for solving large-scale rational eigenvalue problems that are represented as the sum of their polynomial and strictly proper parts as in (2.5). The first key idea is that R-CORK solves the generalized eigenvalue problem associated to the Frobenius companion-like linearization (2.6) previously introduced in [26]. The second key idea is that R-CORK is a structured version of the classical rational Krylov method for solving generalized eigenvalue problems that takes advantage of the particular structure of (2.6). This structure allows us to represent the orthonormal bases of the rational Krylov subspaces of (2.6) in a compact form involving less parameters than the bases of rational Krylov subspaces of the same dimension corresponding to unstructured generalized eigenvalue problems of the same size as the considered linearization. In addition, this compact form can be efficiently and stably updated in each rational Krylov iteration by the use of two levels of orthogonalization in the spirit of the TOAR [27, 13] and the CORK [28] methods for large-scale polynomial eigenvalue problems.


The combined use of the compact representation of rational Krylov subspaces and the two levels of orthogonalization in R-CORK reduces significantly the orthogonalization and the memory costs with respect to a direct application of the classical rational Krylov method to the linearization (2.6). If $n \times n$ is the size of the rational eigenvalue problem, j is the maximum dimension of the considered Krylov subspaces of the linearization, d is the degree of the matrix polynomial $P(\lambda)$ in (2.5), and $s \times s$ is the size of the pencil $(C - \lambda D)$ appearing in (2.5) (note that if $s \leq n$, then s is essentially the rank of the strictly proper part of the rational matrix $R(\lambda)$), then the reduction in costs of R-CORK is appreciable whenever $jd \ll n$ and very considerable if, moreover, $s \ll n$ and $d < j$. In this situation, after j iterations, the orthogonalization cost of R-CORK is $\mathcal{O}(j^2n)$, while the one of classical rational Krylov is $\mathcal{O}(j^2nd)$, and the memory cost of R-CORK is approximately nj numbers, while the one of classical rational Krylov is ndj . These reductions can be combined with an structured implementation of a Krylov-Schur implicit restarting adapted to the compact representation used by R-CORK, which allows us to keep the dimension of the Krylov subspaces moderate without essentially increasing the number of iterations until convergence. The performed numerical experiments confirm all these good properties of R-CORK.

Since many linearizations of rational matrices different from (2.6) have been developed very recently [1, 3] and some of them include the option of considering that the matrix polynomial $P(\lambda)$ in (2.5) is expressed in non-monomial bases, an interesting line of future research on the numerical solution of rational eigenvalue problems is to investigate the potential extension of the R-CORK strategy to other linearizations.


Acknowledgements. The authors sincerely thank Roel Van Beeumen and Karl Meerbergen for answering patiently and very carefully many questions on the CORK method they developed in [28]. Their help has been very important for improving some parts of this manuscript.




(a) Eigenvalues.




(b) Convergence history without restart.




(c) Convergence history with restart.




(d) Dimension of the subspace without restart.



(e) Dimension of the subspace with restart.



(f) Memory cost without restart.



(g) Memory cost with restart.

Figure 5.2: Numerical experiment 5.2

References

- [1] R. Alam and N. Behera. Linearizations for rational matrix functions and Rosenbrock system polynomials. *SIAM J. Matrix Anal. Appl.*, 37(1):354–380, 2016.
- [2] A. Amiraslani, R. M. Corless, and P. Lancaster. Linearization of matrix polynomials expressed in polynomial bases. *IMA J. Numer. Anal.*, 29:141–157, 2009.
- [3] A. Amparan, F. M. Dopico, S. Marcaida, and I. Zaballa. Strong linearizations of rational matrices. MIMS EPrint 2016.51, Manchester Institute for Mathematical Sciences, The University of Manchester, UK, 2016.
- [4] Z. Bai and Y. Su. A second-order Arnoldi method for the solution of the quadratic eigenvalue problem. *SIAM J. Matrix Anal. Appl.*, 26:640–659, 2005.
- [5] T. Betcke, N. J. Higham, V. Mehrmann, C. Schröder, and F. Tisseur. NLEVP: a collection of nonlinear eigenvalue problems. *ACM Trans. Math. Software*, 39(2):Art. 7, 28, 2013.
- [6] G. De Samblanx, K. Meerbergen, and A. Bultheel. The implicit application of a rational filter in the RKS method. *BIT*, 37(4):925–947, 1997.
- [7] F. De Terán, F. M. Dopico, and D. S. Mackey. Fiedler companion linearizations and the recovery of minimal indices. *SIAM J. Matrix Anal. Appl.*, 31:2181–2204, 2010.
- [8] I. Gohberg, P. Lancaster, and L. Rodman. *Matrix Polynomials*. Academic Press, New York, 1982.
- [9] G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, MD, fourth edition, 2013.
- [10] R. Horn and C. Johnson. *Matrix Analysis*. Cambridge University Press, Cambridge, 2nd edition, 2013.
- [11] T. Hwang, W. Lin, J. Liu, and W. Wang. Numerical simulation of a three dimensional quantum dot. *J. Comput. Phys.*, 196:208–232, 2004.
- [12] T. Kailath. *Linear Systems*. Prentice-Hall, Inc., Englewood Cliffs, N.J., 1980.
- [13] D. Kressner and J. Román. Memory-efficient Arnoldi algorithms for linearizations of matrix polynomials in Chebyshev basis. *Numer. Linear Algebra Appl.*, 21(4):569–588, 2014.
- [14] D. Lu, Y. Su, and Z. Bai. Stability analysis of the two-level orthogonal Arnoldi procedure. *SIAM J. Matrix Anal. Appl.*, 37(1):192–214, 2016.
- [15] D. S. Mackey, N. Mackey, C. Mehl, and V. Mehrmann. Structured polynomial eigenvalue problems: Good vibrations from good linearizations. *SIAM J. Matrix Anal. Appl.*, 28:1029–1051, 2006.
- [16] D. S. Mackey, N. Mackey, C. Mehl, and V. Mehrmann. Vector spaces of linearizations for matrix polynomials. *SIAM J. Matrix Anal. Appl.*, 28:971–1004, 2006.
- [17] K. Meerbergen. The quadratic Arnoldi method for the solution of the quadratic eigenvalue problem. *SIAM J. Matrix Anal. Appl.*, 30(4):1463–1482, 2008.
- [18] V. Mehrmann and H. Voss. Nonlinear eigenvalue problems: A challenge for modern eigenvalue methods. *GAMM-Reports*, 27:121–152, 2004.

- [19] S. A. Mohammadi and H. Voss. Variational characterization of real eigenvalues in linear viscoelastic oscillators. Technical report, submitted, 2016.
- [20] H. H. Rosenbrock. *State-space and Multivariable Theory*. Thomas Nelson & Sons, London, 1970.
- [21] A. Ruhe. Rational Krylov sequence methods for eigenvalue computation. *Linear Algebra Appl.*, 58:391–405, 1984.
- [22] A. Ruhe. Rational Krylov: A practical algorithm for large sparse nonsymmetric matrix pencils. *SIAM J. Sci. Comput.*, 19:1535–1551, 1998.
- [23] S. Solov’ev. Preconditioned iterative methods for a class of nonlinear eigenvalue problems. *Linear Algebra Appl.*, 415:210–229, 2006.
- [24] G. W. Stewart. *Matrix Algorithms. Vol. II: Eigensystems*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2001.
- [25] G. W. Stewart. A Krylov-Schur algorithm for large eigenproblems. *SIAM J. Matrix Anal. Appl.*, 23(3):601–614, 2001/02.
- [26] Y. Su and Z. Bai. Solving rational eigenvalue problems via linearization. *SIAM J. Matrix Anal. Appl.*, 32(1):201–216, 2011.
- [27] Y. Su, J. Zhang, and Z. Bai. A compact Arnoldi algorithm for polynomial eigenvalue problems. In *Recent Advances in Numerical Methods for Eigenvalue Problems (RANMEP2008)*, January 2008.
- [28] R. Van Beeumen, K. Meerbergen, and W. Michiels. Compact rational Krylov methods for nonlinear eigenvalue problems. *SIAM J. Matrix Anal. Appl.*, 36(2):820–838, 2015.
- [29] H. Voss. A rational spectral problem in fluid-solid vibration. *Electron. Trans. Numer. Anal.*, 16:94–106, 2003.
- [30] H. Voss. Iterative projection methods for computing relevant energy states of a quantum dot. *BIT*, 44:387–401, 2004.