A Periodic Krylov-Schur Algorithm for Large Matrix Products⋆

Daniel Kressner

Department of Mathematics, Bijenička 30, 10000 Zagreb, Croatia.
e-mail: `kressner@math.hr`.

Summary   Stewart's recently introduced Krylov-Schur algorithm is a modification of the implicitly restarted Arnoldi algorithm which employs reordered Schur decompositions to perform restarts and deflations in a numerically reliable manner. This paper describes a variant of the Krylov-Schur algorithm suitable for addressing eigenvalue problems associated with products of large and sparse matrices. It performs restarts and deflations via reordered periodic Schur decompositions and, by taking the product structure into account, it is capable of achieving qualitatively better approximations to eigenvalues of small magnitude.

Key words   eigenvalue problem, matrix product, Arnoldi method

Mathematics Subject Classification (1991): 65F15, 65F50

1 Introduction

The *product eigenvalue problem* consists of computing eigenvalues and invariant subspaces of a matrix product

$$\Pi = A^{(p)} A^{(p-1)} \cdots A^{(1)}, \tag{1}$$

with the matrices $A^{(1)}, \ldots, A^{(p)} \in \mathbb{C}^{n \times n}$. Instances of this problem arise naturally in a variety of applications, including periodic systems [40], queueing network models [9,37], as well as computational

---

methods for analyzing bifurcations and computing Floquet multipliers [27,28]. On the theoretical side, it has recently been demonstrated that the product eigenvalue problem provides a powerful unifying concept for addressing structured eigenvalue problems which involve, e.g., Hamiltonian, symplectic, pseudosymmetric or unitary matrices [42].

In principal, one could apply any general-purpose eigenvalue solver to the explicitly formed matrix Π. For the sake of numerical stability, however, it is important to take the fact that Π is represented as a matrix product into account and develop numerical methods that work directly on the factors $A^{(1)}, \ldots, A^{(p)}$. The *periodic QR algorithm* [5, 19,38] is such a method for products with small- to medium-sized, dense factors. Requiring $\mathcal{O}(n^2 p)$ memory and $\mathcal{O}(n^3 p)$ computational time, it achieves strong backward stability in the sense that the computed eigenvalues and invariant subspaces correspond to a matrix product with slightly perturbed factors.

In this paper, we propose an algorithm for solving product eigenvalue problems with large, sparse factors. It is an extension of the so-called *Krylov-Schur algorithm* [36], Stewart's modification of the implicitly restarted Arnoldi algorithm [25,31,32], to product eigenvalue problems. To demonstrate the usefulness of our newly developed *periodic Krylov-Schur algorithm*, let us consider the following simple example:

$$A^{(1)} = A^{(2)} = A^{(3)} = \mathrm{diag}(1, 10^{-1}, 10^{-2}, 10^{-3}, \ldots, 10^{-50}).$$

We applied the (standard) Krylov-Schur algorithm as well as the periodic Krylov-Schur algorithm to $A^{(3)} A^{(2)} A^{(1)}$ using random starting vectors and the smallest attainable convergence tolerances. The following table displays the number of correct significant decimal digits of the computed seven largest eigenvalues:

Eigenvalue	Krylov-Schur	Periodic Krylov-Schur
1	15	15
10^{-03}	14	14
10^{-06}	10	14
10^{-09}	8	14
10^{-12}	4	13
10^{-15}	1	12
10^{-18}	0	11

It can be seen that the accuracy of eigenvalues computed by the Krylov-Schur algorithm drops rapidly with decreasing magnitude; the eigenvalue 10^{-18} has no accuracy at all. In comparison, the periodic

Krylov-Schur algorithm is much more accurate; it computes each displayed eigenvalue to at least 11 significant digits correctly.

The rest of this paper is organized as follows. In Section 2, we recall some basic properties of the product eigenvalue problem, its relation to block cyclic matrices and the periodic Schur decomposition. The concepts of periodic Arnoldi and Krylov decompositions are introduced in Sections 3 and 4, respectively. Restarting techniques are described in Section 5, while deflations are the matter of Section 6. Finally, in Section 7, the use of the periodic Krylov-Schur algorithm is demonstrated for two application areas.

2 The Product Eigenvalue Problem

In order to describe the periodic Krylov-Schur algorithm, it is convenient to review some basic facts related to the product eigenvalue problem (1). In this respect, the following *periodic Schur decomposition* plays an important role; just as standard and generalized Schur decompositions play important roles in standard and generalized eigenvalue problems [16].

Theorem 1 (Periodic Schur decomposition [5,19]) *Let $A^{(1)}$, $\ldots$, $A^{(p)} \in \mathbb{C}^{n \times n}$. Then there exist unitary matrices $Q^{(1)}, \ldots, Q^{(p)} \in \mathbb{C}^{n \times n}$ so that*

$$
\begin{aligned}
T^{(p)} &= Q^{(1)H} A^{(p)} Q^{(p)}, \\
T^{(p-1)} &= Q^{(p)H} A^{(p-1)} Q^{(p-1)}, \\
&\vdots \\
T^{(1)} &= Q^{(2)H} A^{(1)} Q^{(1)},
\end{aligned}
\tag{2}
$$

are upper triangular matrices.

The periodic Schur decomposition (2) can be written in the more compact form

$$
T^{(l)} = Q^{(l+1)H} A^{(l)} Q^{(l)}, \quad l = 1, \ldots, p,
$$

if we identify the matrix $Q^{(p+1)}$ with $Q^{(1)}$. More generally speaking, we will make use of the following convention:

Throughout this paper we identify $\star^{(l)}$ with $\star^{(l-1 \bmod p)+1}$, where $\star$ can be replaced by any symbol.

Note that (2) yields a Schur decomposition for the matrix product Π:

$$
Q^{(1)H} \Pi Q^{(1)} = T^{(p)} T^{(p-1)} \cdots T^{(1)} = \diagdown.
\tag{3}
$$

Hence, if $t_{ii}^{(l)}$ denotes the ith diagonal element of $T^{(l)}$, then the n eigenvalues of Π are given by the n products $t_{ii}^{(p)} \cdot t_{ii}^{(p-1)} \cdots t_{ii}^{(1)}$, with $i = 1, \ldots, n$. By a suitable reordering of the periodic Schur decomposition, see [17, 19, 22], we can let the eigenvalues of Π appear in any desirable order on the diagonals of $T^{(l)}$.

   The Schur decomposition (3) also implies that the first k columns of $Q^{(1)}$ span an invariant subspace of Π. More generally, it can be easily seen that if we consider all cyclic permutations of this product,

$$\Pi^{(l)} = A^{(p+l-1)} A^{(p+l-2)} \cdots A^{(l)}, \quad l = 1, \ldots, p, \tag{4}$$

then the first k columns of $Q^{(l)}$ form an invariant subspace of $\Pi^{(l)}$ for each $l \in [1, p]$. Note that in applications related to periodic discrete-time systems, all these invariant subspaces, and not only invariant subspaces of Π, are typically of interest, see [33, 34, 40]. There is a close relationship to certain invariant subspaces of the block cyclic matrix

$$\mathcal{A} = \begin{bmatrix} 0 & & & A^{(p)} \\ A^{(1)} & \ddots & & \\ & \ddots & \ddots & \\ & & A^{(p-1)} & 0 \end{bmatrix}. \tag{5}$$

To see this, let us partition

$$Q^{(l)} = \begin{matrix} k & n-k \\ \begin{bmatrix} X^{(l)} & X_\perp^{(l)} \end{bmatrix} \end{matrix}, \quad T^{(l)} = \begin{matrix} & k & n-k \\ \begin{matrix} k \\ n-k \end{matrix} \begin{bmatrix} A_{11}^{(l)} & A_{12}^{(l)} \\ 0 & A_{22}^{(l)} \end{bmatrix} \end{matrix}.$$

By setting

$$X = X^{(1)} \oplus X^{(2)} \oplus \cdots \oplus X^{(p)}, \quad X_\perp = X_\perp^{(1)} \oplus X_\perp^{(2)} \oplus \cdots \oplus X_\perp^{(p)}, \tag{6}$$

where '$\oplus$' denotes the direct sum of two matrices, and

$$\mathcal{A}_{ij} = \begin{bmatrix} 0 & & & A_{ij}^{(p)} \\ A_{ij}^{(1)} & \ddots & & \\ & \ddots & \ddots & \\ & & A_{ij}^{(p-1)} & 0 \end{bmatrix}, \quad i, j = 1, 2, \tag{7}$$

the identities (2) are equivalent to a block Schur decomposition of $\mathcal{A}$:

$$\mathcal{A}[X, X_\perp] = [X, X_\perp] \begin{bmatrix} \mathcal{A}_{11} & \mathcal{A}_{12} \\ 0 & \mathcal{A}_{22} \end{bmatrix}. \tag{8}$$

In particular, $\mathcal{X} = \text{span}(X)$ is an invariant subspace of $\mathcal{A}$ belonging to the eigenvalues of the block cyclic matrix $\mathcal{A}_{11}$. Not every invariant subspace of $\mathcal{A}$ corresponds to such a block cyclic representation $\mathcal{A}_{11}$, but in the context of product eigenvalue problems it is sufficient to consider this type of subspace.

Definition 1 ([26]) *Let $\mathcal{X}$ be a (pk)-dimensional invariant subspace of a block cyclic matrix $\mathcal{A} \in \mathbb{C}^{pn \times pn}$. If there exist matrices $X^{(1)}$, $X^{(2)}$, ..., $X^{(p)} \in \mathbb{C}^{n \times k}$ so that*

$$\mathcal{X} = \text{span}(X^{(1)} \oplus X^{(2)} \oplus \cdots \oplus X^{(p)}),$$

then $\mathcal{X}$ is called a periodic invariant subspace *of $\mathcal{A}$.*

By direct computation, it can be seen that every periodic invariant subspace corresponds to a block cyclic representation. For $k = 1$, this yields the following well-known relationship between the eigenvalues of Π and $\mathcal{A}$.

Corollary 1 *A scalar λ is an eigenvalue of the matrix product Π having the form (1), if and only if $\lambda^{1/p}, \omega\lambda^{1/p}, \ldots, \omega^{p-1}\lambda^{1/p}$, where ω denotes the pth primitive root of unity, are eigenvalues of the block cyclic matrix $\mathcal{A}$ having the form (5).*

Bhatia [4, Sec. VIII.5] calls a p-tuple of the form $\{\alpha, \omega\alpha, \ldots, \omega^{k-1}\alpha\}$ a *p-Carrollian tuple* in honor of the writer and mathematician Lewis Carroll. Corresponding to a block cyclic representation, a periodic invariant subspace always belongs to eigenvalues that form a set of p-Carrollian tuples.

3 Periodic Arnoldi Decompositions

From the described relationships between product eigenvalue problems and block cyclic matrices it can be concluded that any structure-preserving method for computing p-Carrollian eigenvalue tuples and periodic invariant subspaces of a block cyclic matrix is a suitable method for solving the corresponding product eigenvalue problem.

In the following, we will develop such a structure-preserving method based on the *Krylov subspace*

$$\mathcal{K}_k(\mathcal{A}, u_1) = \text{span}\{u_1, \mathcal{A}u_1, \mathcal{A}^2 u_1, \ldots, \mathcal{A}^{k-1}u_1\}$$

for some starting vector $u_1 \in \mathbb{C}^{pn}$. As we are interested in approximating periodic invariant subspaces, which are spanned by block diagonal matrices, it is reasonable to require that $\mathcal{K}_k(\mathcal{A}, u_1)$ itself admits a block diagonal basis. Indeed, it will be shown that if a starting

vector of the form $u_1 = [u_1^{(1)T}, 0, \ldots, 0]^T$, with $u_1^{(1)} \in \mathbb{C}^n$, is used, then one can construct, by a minor modification of the standard Arnoldi method [16], a decomposition of the following type:

$$\mathcal{A} \cdot \left(U_k^{(1)} \oplus U_k^{(2)} \cdots \oplus U_k^{(p)} \right) = \left(U_{k+1}^{(1)} \oplus U_k^{(2)} \cdots \oplus U_k^{(p)} \right) \cdot \hat{\mathcal{H}}_k. \quad (9)$$

Here, all matrices $U_{k+1}^{(1)} = [u_1^{(1)}, \ldots, u_{k+1}^{(1)}]$ and $U_k^{(l)} = [u_1^{(l)}, \ldots, u_k^{(l)}]$, $l = 2, \ldots, p$, are assumed to have orthonormal columns. Moreover, the factor $\hat{\mathcal{H}}_k$ in (9) takes the form

$$\hat{\mathcal{H}}_k = \begin{bmatrix} 0 & & & \hat{H}_k^{(p)} \\ H_k^{(1)} & \ddots & & \\ & \ddots & \ddots & \\ & & H_k^{(p-1)} & 0 \end{bmatrix} = \begin{bmatrix} & & & \diagdown \\ \diagdown & & & \star \\ & \ddots & & \\ & & \diagdown & \end{bmatrix}, \quad (10)$$

i.e., $H_k^{(1)}, \ldots, H_k^{(p-1)} \in \mathbb{C}^{k \times k}$ are upper triangular matrices while the matrix $\hat{H}_k^{(p)} = \begin{bmatrix} H_k^{(p)} \\ h_{k+1,k}^{(p)} e_k^T \end{bmatrix} \in \mathbb{C}^{(k+1) \times k}$ has (rectangular) upper Hessenberg form. Note that we use e_k to denote the kth unit vector.

Any decomposition of the form (9)–(10) will be called a *periodic Arnoldi decomposition* of order k. Similar decompositions have played a role in Krylov subspace methods for solving linear systems with block cyclic matrices, see [7–9,14,15]. For $p = 1$ the periodic Arnoldi decomposition is nothing but a standard Arnoldi decomposition [35].

Lemma 1 *Consider a kth-order periodic Arnoldi decomposition of the form (9)–(10) and assume that the upper triangular matrices $H_k^{(1)}, \ldots, H_k^{(p-1)}$ are invertible. Moreover, assume that $\hat{H}_k^{(p)}$ is in unreduced Hessenberg form, i.e., all its subdiagonal entries are different from zero. Then the columns of $(U_{k+1}^{(1)} \oplus U_k^{(2)} \cdots \oplus U_k^{(p)})$ form a basis for $\mathcal{K}_{pk+1}(\mathcal{A}, [u_1^{(1)T}, 0, \ldots, 0]^T)$.*

Proof This result is proven by transforming the periodic Arnoldi decomposition into an equivalent standard Arnoldi decomposition. For this purpose, we first permute the $(k+1)$th row of $\hat{\mathcal{H}}_k$ to the bottom row:

$$P_1 \hat{\mathcal{H}}_k = \begin{bmatrix} \mathcal{H}_k \\ h_{k+1,k}^{(p)} e_{pk}^T \end{bmatrix},$$

where $P_1 = [e_1, \ldots, e_k, e_{pk+1}, e_{k+1}, \ldots, e_{pk}]$ denotes the corresponding permutation matrix. Note that $\mathcal{H}_k$ takes the same form as $\hat{\mathcal{H}}_k$ in (10), only with the rectangular Hessenberg matrix $\hat{H}_k^{(p)}$ replaced by the square Hessenberg matrix $H_k^{(p)}$. Next, let us consider the $pk \times pk$ perfect shuffle permutation

$$P_2 = [e_1, e_{1+p}, \ldots, e_{1+(k-1)p}, e_2, e_{2+p}, \ldots, e_{2+(k-1)p}, \ldots,$$
$$e_p, e_{p+p}, \ldots, e_{kp}].$$

Applying this permutation to the rows and columns of $\mathcal{H}_k$ turns it into a Hessenberg matrix $\tilde{\mathcal{H}}_k = P_2 \mathcal{H}_k P_2^T$ whose subdiagonal entries are composed of the diagonal entries of $H_k^{(1)}, \ldots, H_k^{(p-1)}$ and the subdiagonal entries of $H_k^{(p)}$, see also [21]. If we let $V_k = (U_k^{(1)} \oplus U_k^{(2)} \cdots \oplus U_k^{(p)})P_2^T$ and $\hat{P}_2 = P_2 \oplus 1$ then (9) yields the decomposition

$$\mathcal{A}V_k = \left(U_{k+1}^{(1)} \oplus U_k^{(2)} \cdots \oplus U_k^{(p)} \right) P_1^T P_1 \hat{\mathcal{H}}_k P_2^T$$

$$= \left[U_k^{(1)} \oplus U_k^{(2)} \cdots \oplus U_k^{(p)}, \begin{bmatrix} u_{k+1}^{(1)} \\ 0 \end{bmatrix} \right] \hat{P}_2^T \hat{P}_2 \begin{bmatrix} \mathcal{H}_k \\ h_{k+1,k}^{(p)} e_{pk}^T \end{bmatrix} P_2^T$$

$$= \left[V_k, \begin{bmatrix} u_{k+1}^{(1)} \\ 0 \end{bmatrix} \right] \begin{bmatrix} \tilde{\mathcal{H}}_k \\ h_{k+1,k}^{(p)} e_{pk}^T \end{bmatrix}.$$

This is a standard (although particularly structured) Arnoldi decomposition. Moreover, the given assumptions guarantee that the Hessenberg factor is unreduced, which implies that the columns of $\left[V_k, \begin{bmatrix} u_{k+1}^{(1)} \\ 0 \end{bmatrix} \right]$ form a basis for $\mathcal{K}_{pk+1}(\mathcal{A}, [u_1^{(1)T}, 0, \ldots, 0]^T)$, see, e.g., [35, Thm. 5.1.1]. This proves the desired result. $\square$

The proof of Lemma 1 provides one way of relating periodic Arnoldi decompositions to standard Arnoldi decompositions. There is an alternative relationship, which will be explained in the following. For this purpose, let $\mathcal{U}_k = U_k^{(1)} \oplus U_k^{(2)} \cdots \oplus U_k^{(p)}$. Then (9) can be rewritten as

$$\mathcal{A}\mathcal{U}_k = \mathcal{U}_k \mathcal{H}_k + h_{k+1,k}^{(p)} \left[u_{k+1}^{(1)T}, 0, \ldots, 0 \right]^T e_{pk}^T,$$

where $\mathcal{H}_k$ is as $\hat{\mathcal{H}}_k$ in (10) but with $\hat{H}_k^{(p)}$ replaced by the square Hessenberg matrix $H_k^{(p)}$. By repeated application of this equality, we obtain

$$\mathcal{A}^p \mathcal{U}_k = \mathcal{U}_k \mathcal{H}_k^p + \sum_{l=0}^{p-1} h_{k+1,k}^{(p)} \mathcal{A}^{p-l-1} \left[u_{k+1}^{(1)T}, 0, \ldots, 0 \right]^T e_{pk}^T \mathcal{H}_k^l. \quad (11)$$

Note that $\mathcal{A}^p$ is a block diagonal matrix with all cyclic permutations of Π on the block diagonal, i.e., $\Pi^{(1)}, \ldots, \Pi^{(p)}$ defined as in (4). The matrix $\mathcal{H}_k^p$ has a similar structure. Rewriting (11) on the coefficient level yields, after some basic algebraic manipulations,

$$\Pi U_k^{(1)} = U_k^{(1)} \Pi_k^{(1)} + \left(h_{k+1,k}^{(p)} \prod_{l=1}^{p-1} h_{kk}^{(l)} \right) u_{k+1}^{(1)} e_k^T \qquad (12)$$

and, for $i > 1$,

$$\Pi^{(i)} U_k^{(i)} = U_k^{(i)} \Pi_k^{(i)} + \left(h_{k+1,k}^{(p)} \prod_{l=1}^{p-i} h_{kk}^{(l)} \right) A^{(i-1)} \cdots A^{(1)} u_{k+1}^{(1)} e_k^T, \quad (13)$$

where $\Pi_k^{(i)} = H_k^{(p+i-1)} H_k^{(p+i-2)} \cdots H_k^{(i)}$. Note that all matrices $\Pi_k^{(i)}$, $i = 1, \ldots, p$, have upper Hessenberg form, with nonzero subdiagonal entries provided that the conditions of Lemma 1 hold. In this case, the relation (12) is an unreduced standard Arnoldi decomposition for $\Pi^{(1)}$ and thus the columns of $U_{k+1}^{(1)}$ form an orthonormal basis for the Krylov subspace $\mathcal{K}_{k+1}(\Pi, u_1^{(1)})$, see, e.g., [35, Thm. 5.1.1]. Similarly, the columns of $U_k^{(2)}, \ldots, U_k^{(p)}$ form orthonormal bases for the Krylov subspaces $\mathcal{K}_k(\Pi^{(2)}, u_1^{(2)}), \ldots, \mathcal{K}_k(\Pi^{(p)}, u_1^{(p)})$, respectively. In other words, a periodic Arnoldi decomposition for $\mathcal{A}$ simultaneously yields Krylov subspaces for all p cyclic permutations of the matrix product Π.

3.1 The periodic Arnoldi method

Algorithm 1, which we call the *periodic Arnoldi method*, produces a periodic Arnoldi decomposition of the form (9)–(10). It is formally and numerically equivalent to the standard Arnoldi method [16] applied to the block cyclic matrix $\mathcal{A}$ using a starting vector of the form $u_1 = [u_1^{(1)T}, 0, \ldots, 0]^T$, with the notable difference that the columns of the produced Krylov basis are sorted in a particular order. Several remarks concerning the implementation of Algorithm 1 are in order:

1. Algorithm 1 can be implemented such that in each outer loop exactly p matrix-vector multiplications, involving each of the matrices $A^{(1)}, \ldots, A^{(p)}$, are needed. It can be expected that the computational cost of Algorithm 1 is dominated by these matrix-vector multiplications making it comparable to the cost of the standard Arnoldi method applied to the matrix product $\Pi =$

Algorithm 1 Periodic Arnoldi method

Input: Matrices $A^{(1)}, \ldots, A^{(p)} \in \mathbb{C}^{n \times n}$, a starting vector $u_1^{(1)} \in \mathbb{C}^n$ with $\|u_1^{(1)}\|_2 = 1$, and an integer $k < n$.

Output: Matrices $U_{k+1}^{(1)} \in \mathbb{C}^{n \times (k+1)}$, $U_k^{(2)}, \ldots, U_k^{(p)} \in \mathbb{C}^{n \times k}$ having orthonormal columns, upper triangular matrices $H_k^{(1)}, \ldots, H_k^{(p-1)} \in \mathbb{C}^{k \times k}$ and an upper Hessenberg matrix $\hat{H}_k^{(p)} \in \mathbb{C}^{(k+1) \times k}$, defining a kth order periodic Arnoldi decomposition (9)–(10).

$H_0^{(1)} \leftarrow [], \ldots, H_0^{(p-1)} \leftarrow [], \hat{H}_0^{(p)} \leftarrow []$
$U_1^{(1)} \leftarrow \left[u_1^{(1)}\right], U_0^{(2)} \leftarrow [], \ldots, U_0^{(p)} \leftarrow []$
for $j \leftarrow 1, 2, \ldots, k$ **do**
    for $l \leftarrow 1, 2, \ldots, p-1$ **do**
      $h_j^{(l)} \leftarrow U_{j-1}^{(l+1)H} A^{(l)} u_j^{(l)}$ % Note that $U_{j-1}^{(l+1)}$ and $h_j^{(l)}$ are empty for $j = 1$.
      $v \leftarrow A^{(l)} u_j^{(l)} - U_{j-1}^{(l+1)} h_j^{(l)}$
      $h_{jj}^{(l)} \leftarrow \|v\|_2$
      $u_j^{(l+1)} \leftarrow v/h_{jj}^{(l)}, U_j^{(l+1)} \leftarrow \left[U_{j-1}^{(l+1)}, u_j^{(l+1)}\right]$
      $H_j^{(l)} \leftarrow \begin{bmatrix} H_{j-1}^{(l)} & h_j^{(l)} \\ 0 & h_{jj}^{(l)} \end{bmatrix}$
    end for
    $h_j^{(p)} \leftarrow U_j^{(1)H} A^{(p)} u_j^{(p)}$
    $v \leftarrow A^{(p)} u_j^{(p)} - U_j^{(1)} h_j^{(p)}$
    $h_{j+1,j}^{(p)} \leftarrow \|v\|_2$
    $u_{j+1}^{(1)} \leftarrow v/h_{j+1,j}^{(p)}, U_{j+1}^{(1)} \leftarrow \left[U_j^{(1)}, u_{j+1}^{(1)}\right]$
    $\hat{H}_j^{(p)} \leftarrow \begin{bmatrix} \hat{H}_{j-1}^{(p)} & h_j^{(p)} \\ 0 & h_{j+1,j}^{(p)} \end{bmatrix}$
end for

$A^{(p)} A^{(p-1)} \cdots A^{(1)}$. (Note that this statement is only true under the assumption that the matrices $A^{(1)}, \ldots, A^{(p)}$ are subsequently applied to a vector whenever the standard Arnoldi method requests a matrix-vector product. If the matrix product or parts of it can be condensed in a cheap manner, the computational cost of the standard Arnoldi method may be considerably lower.)

2. A major drawback of using Algorithm 1 in comparison to the standard Arnoldi method applied to Π is the increase of memory requirements by roughly a factor of p due to the need for storing each basis $U_{k+1}^{(1)}, U_k^{(2)}, \ldots, U_k^{(p)}$ instead of only one $n \times (k+1)$ basis.

3. Algorithm 1 breaks down as soon as it encounters $h_{jj}^{(l)} = 0$, for $1 \leq l \leq p-1$, or $h_{j+1,j}^{(p)} = 0$. In the latter case, an exact periodic invariant subspace has been found and can be deflated, see Section 6. The first case, $h_{jj}^{(l)} = 0$, also admits the deflation of an invariant subspace, but only after a certain deflation procedure has been applied, see Section 6.1. In both cases, the iteration can be continued with a random, unit vector orthogonal to the previously generated vectors.

4. In exact arithmetic, a newly computed vector $u_j^{(l)}$ will be orthogonal to the columns of $U_{j-1}^{(l)}$. In finite-precision arithmetic, however, this property can get severely violated due to numerical instabilities inherited from the Gram-Schmidt orthonormalization process. This makes it sometimes necessary to reorthogonalize $u_j^{(l)}$ against the columns of $U_{j-1}^{(l)}$. It is straightforward to adapt the reorthogonalization procedures used in the standard Arnoldi method, see [13, 41], for this purpose.

5. So far, we do not know how to construct a satisfactory periodic shift-and-invert Arnoldi method as $(\mathcal{A} - \sigma I)^{-1}$ does not preserve the block cyclic structure of $\mathcal{A}$ as long as $\sigma \neq 0$. Alternatively, one could use the product

$$(\mathcal{A}-\sigma^{(1)}I)^{-1} \cdot (\mathcal{A}-\sigma^{(2)}I)^{-1} \cdots (\mathcal{A}-\sigma^{(p)}I)^{-1} = \left(\mathcal{A}^p - \prod_{l=1}^{p} \sigma^{(l)}I \right)^{-1},$$

where $\{\sigma^{(1)}, \ldots, \sigma^{(p)}\}$ is a p-Carrollian tuple. Such an approach, however, is expected to suffer from numerical instabilities similar to those illustrated by the example in the introduction.

Let us emphasize again that Algorithm 1 is equivalent to the standard Arnoldi method. This implies that the standard methods for computing Ritz values, vectors and bases, see e.g. [35], can be applied to extract approximations to eigenvalues and invariant subspaces from the Krylov subspace basis produced by this algorithm, see also Section 6.

3.2 Convergence properties

It can be expected that the convergence of the periodic Arnoldi method to invariant subspaces often follows closely the convergence of the standard Arnoldi method applied to Π. To see this, let $U_k^{(1)} \oplus \cdots \oplus U_k^{(p)}$ denote the basis for the Krylov subspace $\mathcal{K}_{pk}(\mathcal{A}, [u_1^{(1)T}, 0, \ldots, 0]^T)$ obtained by Algorithm 1. Then the decompositions (12)–(13) imply

that the columns of $U_k^{(1)}, \ldots, U_k^{(p)}$ span bases for $\mathcal{K}_k(\Pi^{(1)}, u_1^{(1)}), \ldots,$ $\mathcal{K}_k(\Pi^{(p)}, u_1^{(p)})$, respectively. Let $\mathcal{X} = \operatorname{span}(X^{(1)} \oplus \cdots \oplus X^{(p)})$ be the desired periodic invariant subspace of dimension $k'p \leq kp$ and set $\mathcal{X}^{(l)} = \operatorname{span}(X^{(l)})$ for $l = 1, \ldots, p$. Because of the block diagonal structure of the bases, we obtain

$$\delta\left(\mathcal{X}, \mathcal{K}_{pk}(\mathcal{A}, [u_1^{(1)T}, 0, \ldots, 0]^T)\right) = \max_{l \in \{1, \ldots, p\}} \delta\left(\mathcal{X}^{(l)}, \mathcal{K}_k(\Pi^{(l)}, u_1^{(l)})\right) \tag{14}$$

where δ denotes the sine of the largest canonical angle between two subspaces [16]. Note that all matrix products $\Pi^{(l)}$ have the same eigenvalues but different eigenvectors. Hence, if the convergence of the standard Arnoldi method applied to $\Pi^{(l)}$ is mainly determined by the eigenvalues of $\Pi^{(l)}$, as considered in [30], then (14) implies that pk steps of the periodic Arnoldi method applied to $\mathcal{A}$ converge as fast as k steps of the standard Arnoldi method applied to $\Pi^{(1)} = \Pi$. Otherwise, the convergence of the Arnoldi method may differ significantly for different l, and the convergence of the periodic Arnoldi method is determined by the slowest one. Such a situation may arise if, e.g., Π is a nearly normal matrix while $\Pi^{(2)}, \ldots, \Pi^{(p)}$ are highly nonnormal matrices [3].

4 Periodic Krylov Decompositions

The applicability of the standard Arnoldi method can be considerably enhanced by employing techniques for restarting and deflating Arnoldi decompositions. Sorensen [31] and Lehoucq [24] have developed efficient implicit restarting and deflation techniques, which have been implemented in the widely used ARPACK software package [25]. Concerning their reliability, it was argued in [23,36] that these techniques sometimes suffer from the forward instability of implicit QR iterations [29], upon which they are based. An elegant solution to the difficulties caused by this effect was proposed by Stewart [36]. It consists of relaxing the definition of an Arnoldi decomposition, leading to the notion of *Krylov decompositions*, and using reliable eigenvalue reordering techniques for restarts and deflations.

In this section, we will extend the notion of Krylov decompositions to periodic Arnoldi decompositions and will use this relaxation in the next section to develop reliable restarting and deflating techniques for the periodic Arnoldi method.

Definition 2 *Let $\mathcal{A}$ be a block cyclic matrix of the form (5). The decomposition*

$$\mathcal{A} \cdot (U_k^{(1)} \oplus U_k^{(2)} \cdots \oplus U_k^{(p)}) = (U_{k+1}^{(1)} \oplus U_k^{(2)} \cdots \oplus U_k^{(p)}) \cdot \hat{\mathcal{B}}_k, \quad (15)$$

is called a periodic Krylov decomposition *of order k if each of the matrices $U_{k+1}^{(1)} = [u_1^{(1)}, \ldots, u_{k+1}^{(1)}]$ and $U_k^{(l)} = [u_1^{(l)}, \ldots, u_k^{(l)}]$, $l = 2, \ldots, p$, has orthonormal columns and if the factor $\hat{\mathcal{B}}_k$ takes the form*

$$\hat{\mathcal{B}}_k = \begin{bmatrix} 0 & & & \hat{B}_k^{(p)} \\ B_k^{(1)} & \ddots & & \\ & \ddots & \ddots & \\ & & B_k^{(p-1)} & 0 \end{bmatrix} \qquad (16)$$

for some matrices $B_k^{(1)}, \ldots, B_k^{(p-1)} \in \mathbb{C}^{k \times k}$ and $\hat{B}_k^{(p)} = \begin{bmatrix} B_k^{(p)} \\ b_k^{(p)H} \end{bmatrix} \in \mathbb{C}^{(k+1) \times k}$.

Although Krylov decompositions are seemingly more general than Arnoldi decompositions, it was shown in [36, Theorem 2.2] that any Krylov decomposition is actually equivalent to an Arnoldi decomposition. The following lemma is an extension of this fact to periodic decompositions.

Lemma 2 *Let*

$$\mathcal{A} \cdot (U_k^{(1)} \oplus U_k^{(2)} \cdots \oplus U_k^{(p)}) = (U_{k+1}^{(1)} \oplus U_k^{(2)} \cdots \oplus U_k^{(p)}) \cdot \hat{\mathcal{B}}_k$$

be a periodic Krylov decomposition. Then there exists a periodic Arnoldi decomposition

$$\mathcal{A} \cdot (\tilde{U}_k^{(1)} \oplus \tilde{U}_k^{(2)} \cdots \oplus \tilde{U}_k^{(p)}) = (\tilde{U}_{k+1}^{(1)} \oplus \tilde{U}_k^{(2)} \cdots \oplus \tilde{U}_k^{(p)}) \cdot \hat{\mathcal{H}}_k \quad (17)$$

so that $\mathrm{span}(\tilde{U}_{k+1}^{(1)}) = \mathrm{span}(U_{k+1}^{(1)})$ and $\mathrm{span}(\tilde{U}_k^{(l)}) = \mathrm{span}(U_k^{(l)})$ for $l = 1, \ldots, p$.

Proof The proof of this result proceeds in a similar fashion as the proof of Theorem 2.2 in [36]. First, let us partition

$$\hat{B}_{k+1}^{(p)} = \begin{bmatrix} B_k^{(p)} \\ b_{k+1}^{(p)H} \end{bmatrix}, \quad B_k^{(p)} \in \mathbb{C}^{k \times k}, \quad b_{k+1}^{(p)} \in \mathbb{C}^k,$$

and construct a unitary matrix Q_0, e.g. a Householder matrix [16], such that $Q_0^H b_{k+1}^{(p)} = \overline{h_{k+1,k}} e_k$ for some $h_{k+1,k} \in \mathbb{C}$. Next, we apply Algorithm 2, a rowwise version of the reduction to periodic Hessenberg form [5], to the matrices

$$B_k^{(1)}, \ B_k^{(2)}, \ \ldots, \ B_k^{(p-2)}, \ Q_0^H B_k^{(p-1)}, \ B_k^{(p)} Q_0.$$

It is simple to check that the unitary factor $Q^{(p)}$ returned by this

---

Algorithm 2 Rowwise reduction to periodic Hessenberg form

Input: Matrices $B^{(1)}, \ldots, B^{(p)} \in \mathbb{C}^{k \times k}$.

Output: Unitary matrices $Q^{(1)}, \ldots, Q^{(p)} \in \mathbb{C}^{k \times k}$ such that $Q^{(1)H} B^{(p)} Q^{(p)}$ is upper Hessenberg and $Q^{(2)H} B^{(1)} Q^{(1)}$, ..., $Q^{(p)H} B^{(p-1)} Q^{(p-1)}$ are upper triangular matrices. The matrices $B^{(1)}, \ldots, B^{(p)}$ are overwritten by the reduced matrices.

Remark: House$_j(x)$ denotes a Householder matrix [16] which maps the leading $j-1$ entries of a vector $x \in \mathbb{C}^k$ to zero without affecting its trailing $n-j$ entries.

$Q^{(1)} \leftarrow I_k, \ \ldots, \ Q^{(p)} \leftarrow I_k$
for $j \leftarrow k, k-1, \ldots, 2$ **do**
   for $l \leftarrow p-1, p-2, \ldots, 1$ **do**
      $x \leftarrow B^{(l)H} e_j$
      $B^{(l)} \leftarrow B^{(l)} \cdot \text{House}_j(x)^H$
      $B^{(l-1)} \leftarrow \text{House}_j(x) \cdot B^{(l-1)}$
      $Q^{(l)} \leftarrow Q^{(l)} \cdot \text{House}_j(x)^H$
   end for
   $x \leftarrow B^{(p)H} e_j$
   $B^{(p)} \leftarrow B^{(l)} \cdot \text{House}_{j-1}(x)^H$
   $B^{(p-1)} \leftarrow \text{House}_{j-1}(x) \cdot B^{(p-1)}$
   $Q^{(p)} \leftarrow Q^{(p)} \cdot \text{House}_{j-1}(x)^H$
end for

---

algorithm satisfies $e_k^T Q^{(p)} = e_k^T$. Note that this requirement is not fulfilled by the original, columnwise algorithm for reduction to periodic Hessenberg form. This implies that setting

$$\tilde{U}_k^{(1)} = U_k^{(1)} Q^{(1)}, \ \ldots, \ \tilde{U}_k^{(p-1)} = U_k^{(p-1)} Q^{(p-1)}, \ \tilde{U}_k^{(p)} = U_k^{(p)} Q_0 Q^{(p)},$$

and $\tilde{U}_{k+1}^{(1)} = [\tilde{U}_k^{(1)}, u_{k+1}^{(1)}]$ yields a periodic Arnoldi decomposition (17), where the factor $\hat{\mathcal{H}}_k$ takes the form (10) with the coefficient matrices

$$H_k^{(1)} = Q^{(2)H} B_k^{(1)} Q^{(1)},$$

$$\vdots$$

$$H_k^{(p-2)} = Q^{(p-1)H} B_k^{(p-2)} Q^{(p-2)},$$

$$H_k^{(p-1)} = Q^{(p)H} Q_0^H B_k^{(p-1)} Q^{(p-1)},$$

$$\hat{H}_k^{(p)} = (Q^{(1)} \oplus 1)^H \hat{B}_k^{(p)} Q_0 Q^{(p)},$$

having Hessenberg-triangular form, which completes the proof. $\square$

The original definition of Krylov decompositions admits non-orthogonal bases. This level of generality is not needed for the algorithms considered in this paper, but by a similar argument as in [36] the statement of Lemma 2 can be extended to the case that $U_{k+1}^{(1)}, U_k^{(2)}, \ldots, U_k^{(p)}$ are only required to have linearly independent columns. In analogy to the notation introduced in [36], a decomposition of the form (15)–(16) is called a *periodic Krylov-Schur decomposition* if all coefficient matrices $B_k^{(1)}, \ldots, B_k^{(p)}$ are upper triangular.

5 Implicit Restarting

One of the drawbacks of Algorithm 1 (as of any Arnoldi method for nonsymmetric matrices) is its need for saving complete Krylov subspace bases. Depending on the speed of convergence, these bases may exceed the available computer memory long before the desired invariant subspace is sufficiently well approximated by the Krylov subspace. Implicit restarting resolves this issue by occasionally disposing of information which does not contain good approximations to the invariant subspace of interest.

Given a periodic Krylov decomposition of order m,

$$\mathcal{A} \cdot (U_m^{(1)} \oplus U_m^{(2)} \oplus \cdots \oplus U_m^{(p)}) = (U_{m+1}^{(1)} \oplus U_m^{(2)} \oplus \cdots \oplus U_m^{(p)}) \cdot \hat{\mathcal{B}}_m, \ (18)$$

where $\hat{\mathcal{B}}_m$ is a block cyclic matrix of the form (16), implicit restarting proceeds as follows. First, unitary matrices $Q_1^{(1)}, \ldots, Q_1^{(p)} \in \mathbb{C}^{m \times m}$ are constructed such that $T_m^{(l)} = Q_1^{(l+1)H} B_m^{(l)} Q_1^{(l)}$ is a periodic Schur decomposition. The eigenvalues of the product

$$\Pi_T = T_m^{(p)} T_m^{(p-1)} \cdots T_m^{(1)}$$

correspond to the m Carrollian tuples of Ritz values of the Krylov decomposition (18). Some of these eigenvalues may be of interest and approximate desired eigenvalues of Π, but some may not. Therefore, the next step of implicit restarting consists of reordering the $k < m$ wanted eigenvalues to the top left corner of the triangular matrix Π_T, using reliable reordering methods as described in [5,6,22]. This corresponds to filtering the starting vector of the periodic Arnoldi method with a degree $(m - k)p$ polynomial (with the zeros being the pth roots of the unwanted eigenvalues), as opposed to a degree $m - k$ polynomial used for the standard Arnoldi method.

In effect, reordering yields another set of unitary matrices $Q_2^{(1)}, \ldots, Q_2^{(p)} \in \mathbb{C}^{m \times m}$ so that

$$
\begin{aligned}
\mathcal{A} \cdot &(U_m^{(1)} Q_1^{(1)} Q_2^{(1)} \oplus U_m^{(2)} Q_1^{(2)} Q_2^{(2)} \oplus \cdots \oplus U_m^{(p)} Q_1^{(p)} Q_2^{(p)}) \\
&= ([U_m^{(1)} Q_1^{(1)} Q_2^{(1)}, u_{m+1}^{(1)}] \oplus U_m^{(2)} Q_1^{(2)} Q_2^{(2)} \oplus \cdots \oplus U_m^{(p)} Q_1^{(p)} Q_2^{(p)}) \cdot \hat{\mathcal{T}}_m,
\end{aligned}
$$

where $\hat{\mathcal{T}}_m$ is a block cyclic matrix of the form (16). The coefficients of $\hat{\mathcal{T}}_m$ are given by triangular matrices $T_m^{(l)} = \begin{bmatrix} T_w^{(l)} & \star \\ 0 & T_u^{(l)} \end{bmatrix}$, $l = 1, \ldots, p$, and

$$
\hat{T}_m^{(p)} = \left[\begin{array}{cc|c} T_w^{(p)} & \star \\ 0 & T_u^{(p)} \\ \hline b_w^{(p)T} & \star \end{array} \right].
$$

Here, the $k \times k$ product $T_w^{(p)} T_w^{(p-1)} \cdots T_w^{(1)}$ contains the wanted eigenvalues of Π_T while the $(m-k) \times (m-k)$ product $T_u^{(p)} T_u^{(p-1)} \cdots T_u^{(1)}$ contains the unwanted eigenvalues. Note that these products correspond to the block cyclic matrices

$$
\mathcal{T}_w = \begin{bmatrix} 0 & & & T_w^{(p)} \\ T_w^{(1)} & \ddots & & \\ & \ddots & \ddots & \\ & & T_w^{(p-1)} & 0 \end{bmatrix}, \quad \mathcal{T}_u = \begin{bmatrix} 0 & & & T_u^{(p)} \\ T_u^{(1)} & \ddots & & \\ & \ddots & \ddots & \\ & & T_u^{(p-1)} & 0 \end{bmatrix}. \quad (19)
$$

Finally, the constructed periodic Krylov-Schur decomposition is truncated. By letting $\tilde{U}_k^{(l)}$ contain the first k columns of $U_m^{(l)} Q_1^{(l)} Q_2^{(l)}$, $l = 1, \ldots, p$, and setting $\tilde{u}_{k+1}^{(1)} = u_{m+1}^{(1)}$, we obtain the following Krylov-Schur decomposition of order k:

$$
\mathcal{A} \cdot (\tilde{U}_k^{(1)} \oplus \tilde{U}_k^{(2)} \oplus \cdots \oplus \tilde{U}_k^{(p)}) = ([\tilde{U}_k^{(1)}, \tilde{u}_{k+1}^{(1)}] \oplus \tilde{U}_k^{(2)} \oplus \cdots \oplus \tilde{U}_k^{(p)}) \cdot \hat{\mathcal{T}}_w,
$$

where $\hat{\mathcal{T}}_w$ is almost identical to $\mathcal{T}_w$ in (19), with the only difference being that the $k \times k$ coefficient matrix $T_w^{(p)}$ is replaced by the $(k+1) \times k$ matrix $\hat{T}_w^{(p)} = \begin{bmatrix} T_w^{(p)} \\ b_w^{(p)T} \end{bmatrix}$.

The described procedure of reduction to periodic Krylov-Schur decomposition, reordering and truncation is depicted in Figure 1. Note that the Krylov subspace corresponding to the truncated decomposition is the same subspace that would have been obtained if implicit restarting with the filter polynomial

$$
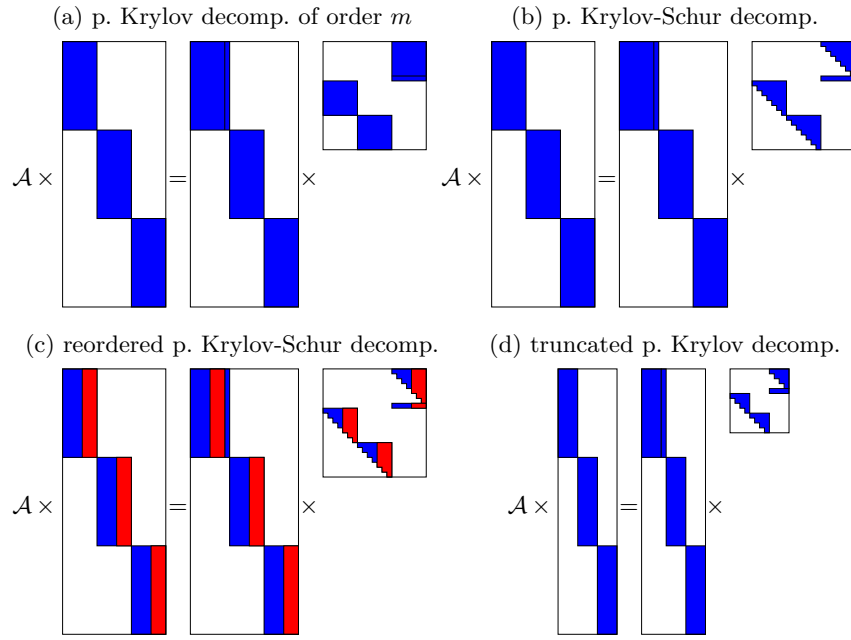\psi(z) = \prod_{\lambda \in \lambda(\mathcal{T}_u)} (z - \lambda),
$$

(a) p. Krylov decomp. of order m (b) p. Krylov-Schur decomp.

(c) reordered p. Krylov-Schur decomp.   (d) truncated p. Krylov decomp.

Fig. 1. Restarting a periodic Krylov decomposition.

see [31], had been applied to the standard Arnoldi decomposition corresponding to (18). The notable difference of our method is that it exploits and preserves the structure of periodic Krylov decompositions.

6 Deflation

After a periodic Krylov decomposition has been truncated to order k, it can again be expanded to a decomposition of order m by applying a variant of the periodic Arnoldi method, Algorithm 1. This process of truncation and expansion is repeated until convergence occurs. We suggest the use of a convergence criterion similar to the one used in [36] for deflating invariant subspaces from Krylov decompositions. In order to preserve the structure of periodic Krylov decompositions, we will only deflate invariant subspaces that belong to a p-Carrollian tuple of Ritz values.

For the purpose of describing this deflation strategy in detail, assume that we have an mth order periodic Krylov decomposition, where $d < m$ p-Carrollian tuples of Ritz values have already been deflated. This corresponds to a decomposition of the following parti-

tioned form:

$$
\begin{aligned}
\mathcal{A} \cdot ([Q_d^{(1)}, U_{m-d}^{(1)}] &\oplus [Q_d^{(2)}, U_{m-d}^{(2)}] \cdots \oplus [Q_d^{(p)}, U_{m-d}^{(p)}]) \\
&= ([Q_d^{(1)}, U_{m-d+1}^{(1)}] \oplus [Q_d^{(2)}, U_{m-d}^{(2)}] \cdots \oplus [Q_d^{(p)}, U_{m-d}^{(p)}]) \cdot \hat{\mathcal{B}}_m,
\end{aligned}
\tag{20}
$$

where $Q_d^{(1)}, \ldots, Q_d^{(p)} \in \mathbb{C}^{n \times d}$ and $\hat{\mathcal{B}}_m$ is a block cyclic matrix of the form (16) with accordingly partitioned coefficient matrices

$$
B_m^{(l)} = \begin{bmatrix} T_d^{(l)} & \star \\ 0 & B_{m-d}^{(l)} \end{bmatrix}, \; l = 1, \ldots, p-1, \quad \hat{B}_m^{(p)} = \begin{bmatrix} T_d^{(p)} & \star \\ 0 & \hat{B}_{m-d}^{(p)} \end{bmatrix}.
\tag{21}
$$

Let $\mathcal{T}_d$ be the block cyclic matrix associated with the (usually upper triangular) matrices $T_d^{(1)}, \ldots, T_d^{(p)} \in \mathbb{C}^{d \times d}$. The structure of (20)–(21) implies that the columns of the matrix $Q_d = Q_d^{(1)} \oplus \cdots \oplus Q_d^{(p)}$ span a periodic invariant subspace belonging to the eigenvalues of $\mathcal{T}_d$. In other words, the eigenvalues of $\mathcal{T}_d$ have already been deflated.

To deflate another p-Carrollian tuple of Ritz values, we assume, e.g. by using a periodic Schur decomposition, that the submatrices $B_{m-d}^{(1)}, \ldots, B_{m-d}^{(p-1)}, \hat{B}_{m-d}^{(p)}$ in (21) take the form

$$
B_m^{(l)} = \begin{bmatrix} \lambda^{(l)} & \star \\ 0 & \tilde{B}_{m-d-1}^{(l)} \end{bmatrix}, \; l = 1, \ldots, p-1, \quad \hat{B}_m^{(p)} = \begin{bmatrix} \lambda^{(p)} & \star \\ 0 & \tilde{B}_{m-d-1}^{(p)} \\ \hline \tilde{b}_1^{(p)} & \tilde{b}_{m-d-1}^{(p)H} \end{bmatrix}
$$

for some scalars $\lambda^{(1)}, \ldots, \lambda^{(p)}$. The eigenvalues of the cyclic matrix

$$
\Lambda = \begin{bmatrix} 0 & & & \lambda^{(p)} \\ \lambda^{(1)} & \ddots & & \\ & \ddots & \ddots & \\ & & \lambda^{(p-1)} & 0 \end{bmatrix}
$$

are regarded as a converged p-Carrollian tuple of Ritz values if the scalar $\tilde{b}_1^{(p)}$ satisfies

$$
|\tilde{b}_1^{(p)}| \leq \max\{\mathbf{u} \cdot \|\hat{B}_m^{(p)}\|_F, \texttt{tol} \cdot |\lambda^{(p)}|\},
\tag{22}
$$

where $\mathbf{u}$ denotes the machine precision and $\texttt{tol} > 0$ is a chosen user tolerance. Being in the spirit of standard convergence criteria for Arnoldi methods [25, Sec. 4.6], this criterion implies that the $n \times p$ matrix $V = U_{m-d}^{(1)} e_1 \oplus \cdots \oplus U_{m-d}^{(p)} e_1$ satisfies

$$
\|(I - Q_d Q_d^H) \mathcal{A}(I - Q_d Q_d^H) V - V \Lambda\|_F \leq \max\{\mathbf{u} \cdot \|\hat{B}_m^{(p)}\|_F, \texttt{tol} \cdot |\lambda^{(p)}|\}.
$$

Thus, the columns of V nearly span an invariant subspace belonging to the eigenvalues of Λ, in the sense that the residual of the pair (V, Λ) is satisfactorily small. If the considered p-Carrollian tuple does not satisfy (22), we may test any other tuple of Ritz values by reordering the periodic Schur form of the matrix product

$$\begin{bmatrix} \lambda^{(p)} & \star \\ 0 & \tilde{B}_{m-d-1}^{(p)} \end{bmatrix} \begin{bmatrix} \lambda^{(p-1)} & \star \\ 0 & \tilde{B}_{m-d-1}^{(p-1)} \end{bmatrix} \cdots \begin{bmatrix} \lambda^{(1)} & \star \\ 0 & \tilde{B}_{m-d-1}^{(1)} \end{bmatrix}. \qquad (23)$$

As already pointed out in [36], the use of a convergence criterion of the form (22) leads to a numerically backward stable method for approximating invariant subspaces. To see this for the periodic Krylov-Schur algorithm, let us assume that the d deflated p-Carrollian tuples of Ritz values in (20) have been deflated based upon a criterion of the form (22). Then the matrix $Q_d = (Q_d^{(1)} \oplus \cdots \oplus Q_d^{(p)})$ satisfies

$$\mathcal{A}Q_d - Q_d \mathcal{T}_d = R_d,$$

where $R_d \in \mathbb{C}^{np \times dp}$ is a block cyclic matrix with

$$\|R_d\|_F \leq \sqrt{d} \max\{\mathbf{u} \cdot \|A^{(p)}\|_F, \texttt{tol} \cdot \|T_d^{(p)}\|_F\}.$$

The subspace spanned by the columns of Q_d is the *exact* periodic invariant subspace of the slightly perturbed block cyclic matrix $\mathcal{A} - R_d Q_d^H$. Neglecting the effects of roundoff errors, this shows the strong backward stability [11] of the periodic Krylov-Schur algorithm for computing periodic invariant subspaces of block cyclic matrices. Similarly as for the standard Krylov-Schur algorithm [36], it can be shown that roundoff errors do not introduce numerical instabilities, if Algorithm 1 is carefully implemented.

Remark 1 In the case of real factors $A^{(1)}, \ldots, A^{(p)}$, the realness of periodic Krylov decompositions can be preserved by employing real periodic Schur decompositions [5] and reordering/deflating complex conjugate pairs of eigenvalues together.

6.1 Deflation of singular factors

In the case that one of the triangular factors $H_m^{(1)}, \ldots, H_m^{(p-1)}$ in an mth order periodic Arnoldi decomposition (9)–(10) happens to be (almost) singular, then the assumption of Lemma 1 is not satisfied and the basis produced by Algorithm 1 does not span a Krylov subspace. Such a situation arises, e.g., if Algorithm 1 encounters a zero

vector v in the inner loop and continues the iteration with a random, unit vector orthogonal to the previously generated vectors in $U_{j-1}^{(l+1)}$.

In the following, it will be shown that a zero diagonal element in the triangular factors admits the deflation of an invariant subspace by using a deflation procedure similar to the one used in the periodic QR algorithm for dealing with singular triangular factors, see [5,38]. Let us outline this deflation for a fifth order periodic Arnoldi decomposition with $p = 3$, where the third diagonal element of the triangular matrix $H_5^{(2)}$ happens to be zero:

$$
\hat{H}_5^{(3)} H_5^{(2)} H_5^{(1)} =
\begin{bmatrix}
x & x & x & x & x \\
x & x & x & x & x \\
0 & x & x & x & x \\
0 & 0 & x & x & x \\
0 & 0 & 0 & x & x \\
0 & 0 & 0 & 0 & x
\end{bmatrix}
\begin{bmatrix}
x & x & x & x & x \\
0 & x & x & x & x \\
0 & 0 & 0 & x & x \\
0 & 0 & 0 & x & x \\
0 & 0 & 0 & 0 & x
\end{bmatrix}
\begin{bmatrix}
x & x & x & x & x \\
0 & x & x & x & x \\
0 & 0 & x & x & x \\
0 & 0 & 0 & x & x \\
0 & 0 & 0 & 0 & x
\end{bmatrix}.
$$

Applying a sequence of two Givens rotations [16], the subdiagonal elements $(2,1)$ and $(3,2)$ of $\hat{H}_5^{(3)}$ can be annihilated at the expense of introducing two nonzeros in the subdiagonal of $H_5^{(1)}$:

$$
\begin{bmatrix}
\hat{x} & \hat{x} & \hat{x} & \hat{x} & \hat{x} \\
\hat{0} & \hat{x} & \hat{x} & \hat{x} & \hat{x} \\
0 & \hat{0} & \hat{x} & \hat{x} & \hat{x} \\
0 & 0 & x & x & x \\
0 & 0 & 0 & x & x \\
0 & 0 & 0 & 0 & x
\end{bmatrix}
\begin{bmatrix}
x & x & x & x & x \\
0 & x & x & x & x \\
0 & 0 & 0 & x & x \\
0 & 0 & 0 & x & x \\
0 & 0 & 0 & 0 & x
\end{bmatrix}
\begin{bmatrix}
\hat{x} & \hat{x} & \hat{x} & x & x \\
\hat{x} & \hat{x} & \hat{x} & x & x \\
0 & \hat{x} & \hat{x} & x & x \\
0 & 0 & 0 & x & x \\
0 & 0 & 0 & 0 & x
\end{bmatrix}.
$$

These two newly created nonzeros can be annihilated in a similar fashion, but now – owing to the zero diagonal element – only one nonzero entry is introduced in $H_5^{(2)}$:

$$
\begin{bmatrix}
x & x & x & x & x \\
0 & x & x & x & x \\
0 & 0 & x & x & x \\
0 & 0 & x & x & x \\
0 & 0 & 0 & x & x \\
0 & 0 & 0 & 0 & x
\end{bmatrix}
\begin{bmatrix}
\hat{x} & \hat{x} & \hat{x} & x & x \\
\hat{x} & \hat{x} & \hat{x} & x & x \\
0 & 0 & 0 & x & x \\
0 & 0 & 0 & x & x \\
0 & 0 & 0 & 0 & x
\end{bmatrix}
\begin{bmatrix}
\hat{x} & \hat{x} & \hat{x} & \hat{x} & \hat{x} \\
\hat{0} & \hat{x} & \hat{x} & \hat{x} & \hat{x} \\
0 & \hat{0} & \hat{x} & \hat{x} & \hat{x} \\
0 & 0 & 0 & x & x \\
0 & 0 & 0 & 0 & x
\end{bmatrix}.
$$

Annihilating this entry returns the product to periodic Hessenberg form:

$$
\begin{bmatrix}
\hat{x} & \hat{x} & x & x & x \\
\hat{x} & \hat{x} & x & x & x \\
0 & 0 & x & x & x \\
0 & 0 & x & x & x \\
0 & 0 & 0 & x & x \\
0 & 0 & 0 & 0 & x
\end{bmatrix}
\begin{bmatrix}
\hat{x} & \hat{x} & \hat{x} & \hat{x} & \hat{x} \\
\hat{0} & \hat{x} & \hat{x} & \hat{x} & \hat{x} \\
0 & 0 & 0 & x & x \\
0 & 0 & 0 & x & x \\
0 & 0 & 0 & 0 & x
\end{bmatrix}
\begin{bmatrix}
x & x & x & x & x \\
0 & x & x & x & x \\
0 & 0 & x & x & x \\
0 & 0 & 0 & x & x \\
0 & 0 & 0 & 0 & x
\end{bmatrix}. \qquad (24)
$$

This yields a deflated periodic invariant subspace of dimension $2p$. Although another p-Carrollian tuple of eigenvalues is known to consist of zero eigenvalues there seems to be no straightforward way to deflate the corresponding periodic invariant subspace unless the $(4,3)$ subdiagonal entry of $\hat{H}_5^{(3)}$ is sufficiently small. A simple remedy is to ignore this additional information and perform an explicit restart with a random starting vector orthogonal to the deflated subspace.

A technically more involved remedy is to allow varying dimensions and relabel the periodic Krylov decomposition. Let us repartition (24) as follows:

$$
\begin{bmatrix}
x & x & x & x & x \\
x & x & x & x & x \\
0 & 0 & x & x & x \\
0 & 0 & x & x & x \\
0 & 0 & 0 & x & x \\
0 & 0 & 0 & 0 & x
\end{bmatrix}
\begin{bmatrix}
x & x & x & x & x \\
0 & x & x & x & x \\
0 & 0 & 0 & x & x \\
0 & 0 & 0 & x & x \\
0 & 0 & 0 & 0 & x
\end{bmatrix}
\begin{bmatrix}
x & x & x & x & x \\
0 & x & x & x & x \\
0 & 0 & x & x & x \\
0 & 0 & 0 & x & x \\
0 & 0 & 0 & 0 & x
\end{bmatrix}.
$$

By applying two further steps of the periodic Arnoldi method, Algorithm 1, one obtains the following periodic Hessenberg form:

$$
\begin{bmatrix}
x & x & x & x & x \\
x & x & x & x & x \\
0 & 0 & x & x & x \\
0 & 0 & x & x & x \\
0 & 0 & 0 & x & x \\
0 & 0 & 0 & 0 & x
\end{bmatrix}
\begin{bmatrix}
x & x & x & x & x & x \\
0 & x & x & x & x & x \\
0 & 0 & 0 & x & x & x \\
0 & 0 & 0 & x & x & x \\
0 & 0 & 0 & 0 & x & x \\
0 & 0 & 0 & 0 & 0 & x
\end{bmatrix}
\begin{bmatrix}
x & x & x & x & x & x \\
0 & x & x & x & x & x \\
0 & 0 & x & x & x & x \\
0 & 0 & 0 & x & x & x \\
0 & 0 & 0 & 0 & x & x \\
0 & 0 & 0 & 0 & 0 & x
\end{bmatrix}.
$$

This corresponds to a periodic Arnoldi decomposition of the form

$$
A^{(3)} \cdot \overset{2\qquad 3}{\left[\, Q^{(3)}, U^{(3)} \,\right]} = \overset{3\qquad 3}{\left[\, Q^{(1)}, U^{(1)} \,\right]} \cdot \begin{smallmatrix} 3 \\ 3 \end{smallmatrix} \overset{2\qquad\quad 3}{\begin{bmatrix} T^{(3)} & \star \\ 0 & H^{(3)} \end{bmatrix}},
$$

$$
A^{(2)} \cdot \overset{3\qquad 3}{\left[\, Q^{(2)}, U^{(2)} \,\right]} = \overset{2\qquad 3\qquad 1}{\left[\, Q^{(3)}, U^{(3)}, u^{(3)} \,\right]} \cdot \begin{smallmatrix} 2 \\ 4 \end{smallmatrix} \overset{3\qquad\quad 3}{\begin{bmatrix} T^{(2)} & \star \\ 0 & \hat{H}^{(2)} \end{bmatrix}}, (25)
$$

$$A^{(1)} \cdot \left[\, \overset{3}{Q^{(1)}}, \overset{3}{U^{(1)}} \,\right] = \left[\, \overset{3}{Q^{(2)}}, \overset{3}{U^{(2)}} \,\right] \cdot \begin{matrix} 3 \\ 3 \end{matrix} \begin{bmatrix} \overset{3}{T^{(1)}} & \overset{3}{\star} \\ 0 & H^{(1)} \end{bmatrix}.$$

Note that the deflated matrix product $T^{(3)}T^{(2)}T^{(1)}$ consists of rectangular factors, which accounts for the discovered zero eigenvalue. A slight variation of the periodic QR algorithm can be used to compute eigenvalues and invariant subspaces of such matrix products in a numerically backward stable manner, see [39] for more details. Also, the subspace $\mathcal{X}$ spanned by the columns of $Q^{(1)} \oplus Q^{(2)} \oplus Q^{(3)}$ is an invariant subspace of $\mathcal{A}$. Strictly speaking, $\mathcal{X}$ is not a periodic invariant subspace in the sense of Definition 1, as the diagonal blocks in the basis of $\mathcal{X}$ do not have the same number of columns. Nevertheless, the columns of $Q^{(1)}$, $Q^{(2)}$, and $Q^{(3)}$ span invariant subspaces of the matrix products $\Pi^{(1)}, \Pi^{(2)}$, and $\Pi^{(3)}$, respectively.

The undeflated part in (25) corresponds to the decomposition

$$A^{(3)}U^{(3)} = U^{(1)}H^{(3)}, \; A^{(2)}U^{(2)} = [U^{(3)}, u^{(3)}]\hat{H}^{(2)}, \; A^{(1)}U^{(1)} = U^{(2)}H^{(1)}$$

By the index transformation

$$(2) \to (3), \quad (3) \to (1), \quad (1) \to (2),$$

it can be seen that this decomposition is actually a periodic Arnoldi decomposition. Hence, the periodic Arnoldi method can be continued from this decomposition.

7 Numerical Examples

In this section, we demonstrate the use of the periodic Krylov-Schur algorithm for two application areas: model reduction and periodic discrete-time systems.

Before, we consider the simple example from the introduction:

$$A^{(1)} = A^{(2)} = A^{(3)} = \mathrm{diag}(1, 10^{-1}, 10^{-2}, 10^{-3}, \ldots, 10^{-50}). \quad (26)$$

Figure 2 displays the approximation error

$$\min\{|\lambda_i - \sigma_j| : \sigma_j \text{ is Ritz value}\} \quad (27)$$

for each of the 7 largest eigenvalues $1, 10^{-3}, \ldots, 10^{-18}$ versus the number of (product) Arnoldi steps. It can be seen that these errors stagnate at a value above 10^{-20} for the Krylov-Schur algorithm, while the product Krylov-Schur algorithm is capable of computing much more accurate approximations to the smaller eigenvalues (which correspond to the lower error curves). All computations in this and the following numerical examples have been performed in MATLAB 6.5.
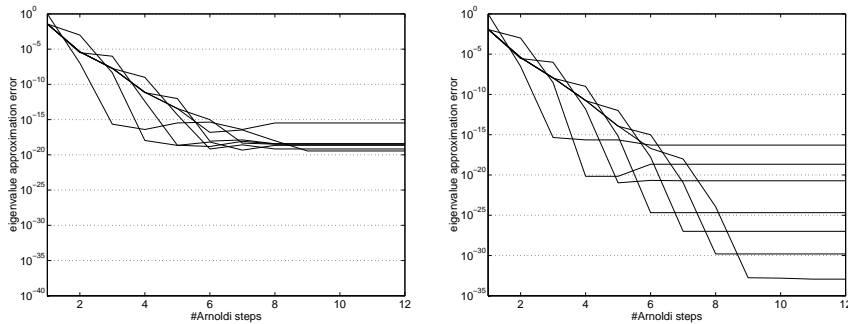
Fig. 2. Approximation errors of the 7 largest eigenvalues for the standard (left plot) and the periodic (right plot) Krylov-Schur algorithm applied to the product $A^{(3)}A^{(2)}A^{(1)}$ as defined in (26).

7.1 Model reduction and Hankel singular values

Let us consider a continuous, linear time-invariant (LTI) system

$$
\begin{aligned}
\dot{x}(t) &= Ax(t) + Bu(t), \\
y(t) &= Cx(t),
\end{aligned} \tag{28}
$$

where $x(t) \in \mathbb{C}^n$, $u(t) \in \mathbb{C}^m$, and $y(t) \in \mathbb{C}^q$ represent the state, input, and output vectors at time t, respectively, while the system matrices $A \in \mathbb{C}^{n\times n}, B \in \mathbb{C}^{n\times m}, C \in \mathbb{C}^{q\times n}$ are assumed to be constant. Model reduction aims at approximating (28) by a reduced LTI system with a similar input-output-behavior but with a much smaller number $\tilde{n} \ll n$ of states.

Some popular model reduction techniques, such as balanced truncation, are closely related to the two Lyapunov equations

$$
AP + PA^H = -BB^H, \qquad A^HQ + QA = -C^HC, \tag{29}
$$

see [1]. It will be assumed that the system (28) is stable, i.e., all eigenvalues of A lie in the open left half plane. In this case, one can directly compute Cholesky factors R and S such that $P = S^HS$ and $Q = R^HR$ are the unique solutions of (29), see, e.g., [18,20]. Balanced truncation proceeds by computing a singular value decomposition

$$
RS^H = U\Sigma V^H, \tag{30}
$$

where the matrices U, V are unitary and Σ is a diagonal matrix with the so-called *Hankel singular values* on its diagonal. This results in a transformation matrix $T = S^HV\Sigma^{-1/2}$ which in turn yields a balanced system with system matrices $T^{-1}AT$, $T^{-1}B$, TC. This balanced system is reduced by maintaining only those parts that belong
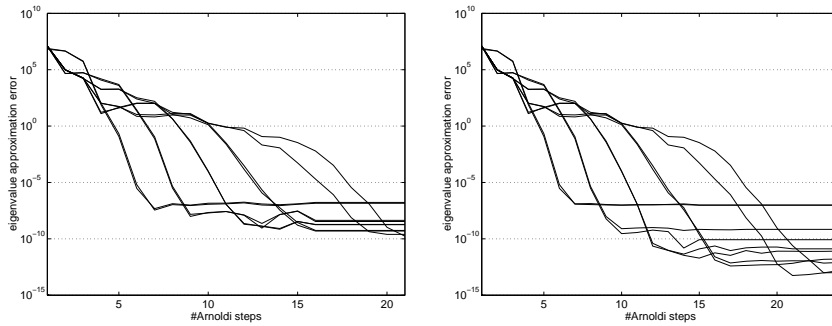
Fig. 3. Approximation errors of the 10 largest eigenvalues for the standard (left plot) and the periodic (right plot) Krylov-Schur algorithm applied to the matrix product Π corresponding to a discretized clamped beam model.

to non-negligible Hankel singular values. For this purpose, it is important to identify these values correctly.

Computing the Hankel singular values is equivalent to computing the positive square roots of the eigenvalues of the matrix product $\Pi = SR^H RS^H$. In the following, we consider a practical example to compare the standard Krylov-Schur algorithm with the periodic Krylov-Schur algorithm applied to Π. For this purpose, we have used the discretized model of a clamped beam as described in [2,12], for which $n = 348$ and $m = q = 1$. The corresponding Cholesky factors R and S as well as the Hankel singular values can be obtained from the model reduction benchmark collection [12]. Figure 3 displays the eigenvalue approximation errors, see (27), for the 10 largest eigenvalues of Π:

$$\{5.7 \cdot 10^6, 4.7 \cdot 10^6, 7.4 \cdot 10^4, 7.1 \cdot 10^4, 2.0 \cdot 10^3, 1.9 \cdot 10^3, 1.1 \cdot 10^2, 13, 9.7\}$$

(only the leading two significant digits are displayed). A restart was applied after 17 Arnoldi steps, truncating the (periodic) Arnoldi decomposition to a (periodic) Krylov-Schur decomposition which maintains the 12 largest Ritz values. The errors in the eigenvalue approximations produced by the Krylov-Schur algorithm stagnate at a level above 10^{-10}. Again, the periodic Krylov-Schur algorithm produces more accurate approximations; the lower error curves in the right plot of Figure 3 correspond to the smaller eigenvalues.

7.2 Periodic discrete-time systems

A periodic discrete-time system takes the form

$$x_{l+1} = A^{(l)} x_l + B^{(l)} u_l,$$

$$y_{l+1} = C^{(l)}x_l, \qquad l = 0, 1, \ldots,$$

where $x_l \in \mathbb{C}^n, u_l \in \mathbb{C}^m$, and $y_l \in \mathbb{C}^p$ are the state, input, and output vectors, respectively. The system matrices $A^{(l)}, B^{(l)}, C^{(l)}$ have appropriate dimensions and are assumed to be periodic for some fixed period $p \geq 1$, i.e., $A^{(l+p)} = A^{(l)}, B^{(l+p)} = B^{(l)}, C^{(l+p)} = C^{(l)}$. The Floquet multipliers of this system, which determine the asymptotic behavior of x_l for $l \to \infty$, are given by the pth roots of the eigenvalues of the matrix product $A^{(p-1)}A^{(p-2)}\cdots A^{(0)}$.

In view of the lack of publicly available large-scale periodic discrete-time systems, we consider a system of state dimension $n = 12$ that models the dynamics of nitrogen absorption, distribution, and translocation in citrus trees [10]. This system has period $p = 365$ and its nonzero Floquet multipliers are the 365th roots of

$$0.1943, \ 0.1614, \ 9.834 \cdot 10^{-3}, \ 1.048 \cdot 10^{-3}, \ 1.013 \cdot 10^{-14}$$

(only the leading four significant digits are displayed). To compute these values, we applied the (periodic) Krylov-Schur algorithm to $\Pi = A^{(364)}A^{(363)}\cdots A^{(0)}$. The parameters were set such that 5 eigenvalues are to be computed and restarts are performed if the dimension of the Krylov subspace exceeds $9p + 1$, in which case the subspace is truncated to dimension $6p + 1$ (with $p = 1$ for the standard Krylov-Schur algorithm). We observed that the standard Krylov-Schur algorithm was not capable of computing any reasonable approximation to the smallest nonzero eigenvalue ($1.013 \cdot 10^{-14}$), instead $9.471 \cdot 10^{-11} + 9.952 \cdot 10^{-11}\imath$ was returned as the fifth-largest converged Ritz value. The periodic Krylov-Schur algorithm, however, computed the leading five significant decimal digits of this eigenvalue correctly.

It is worth mentioning again that many computational tasks related to periodic discrete-time systems require the complete information contained in the (approximate) periodic invariant subspaces of Π [33, 34, 40]. While being a natural by-product of the periodic Krylov-Schur algorithm, computing this information reliably by the standard Krylov-Schur algorithm is an expensive and difficult task.

8 Conclusions and Future Work

We have introduced the periodic Krylov-Schur algorithm, an Arnoldi-like algorithm tailored to matrix products. In terms of accuracy, it has been demonstrated that this algorithm compares favorably with the standard Krylov-Schur algorithm. However, several important issues such as reliable shift-and-invert strategies, harmonic Ritz values, refined Ritz vectors and null-space purification remain to be addressed

in order to make the periodic Krylov-Schur algorithm as robust and practicable as, e.g., ARPACK [25].

9 Final Remark and Acknowledgments

References

1. A. C. Antoulas. *Approximation of Large-Scale Dynamical Systems*. SIAM, Philadelphia, PA, 2005.
2. A. C. Antoulas, D. C. Sorensen, and S. Gugercin. A survey of model reduction methods for large-scale systems. In *Structured matrices in mathematics, computer science, and engineering, I (Boulder, CO, 1999)*, volume 280 of *Contemp. Math.*, pages 193–219. Amer. Math. Soc., Providence, RI, 2001.
3. C. A. Beattie, M. Embree, and D. C. Sorensen. Convergence of polynomial restart Krylov methods for eigenvalue computations. *SIAM Rev.*, 47(3):492–515, 2005.
4. R. Bhatia. *Matrix Analysis*. Springer-Verlag, New York, 1997.
5. A. Bojanczyk, G. H. Golub, and P. Van Dooren. The periodic Schur decomposition; algorithm and applications. In *Proc. SPIE Conference*, volume 1770, pages 31–42, 1992.
6. A. Bojanczyk and P. Van Dooren. On propagating orthogonal transformations in a product of 2×2 triangular matrices. In L. Reichel, A. Ruttan, and R. S. Varga, editors, *Numerical Linear Algebra (Kent, OH, 1992)*, pages 1–9. de Gruyter, 1993.
7. W. Bomhof. *Iterative and parallel methods for linear systems, with applications in circuit simulation*. PhD thesis, Universiteit Utrecht, Faculteit der Wiskunde en Informatica, Utrecht, Netherlands, 2001.
8. W. Bomhof and H. van der Vorst. A parallelizable GMRES-type method for p-cyclic matrices, with applications in circuit simulation. In U. Van Rienen, M. Gunther, and D. Hecht, editors, *Scientific computing in electrical engineering: proceedings of the 3rd international workshop, August 20–23, 2000, Warnemünde, Germany*, volume 18 of *Lecture Notes in Computational Science and Engineering*, pages 293–300, Berlin, Heidelberg, London, 2001. Springer Verlag.
9. F. Bonhoure, Y. Dallery, and W. J. Stewart. On the use of periodicity properties for the efficient numerical solution of certain Markov chains. *Numer. Linear Algebra Appl.*, 1(3):265–286, 1994.
10. R. Bru, R. Cantó, and B. Ricarte. Modelling nitrogen dynamics in citrus trees. *Mathematical and Computer Modelling*, 38:975–987, 2003.

11. J. R. Bunch. The weak and strong stability of algorithms in numerical linear algebra. *Linear Algebra Appl.*, 88/89:49–66, 1987.

12. Y. Chahlaoui and P. Van Dooren. Benchmark examples for model reduction of linear time-invariant dynamical systems. In P. Benner, V. Mehrmann, and D. C. Sorensen, editors, *Dimension Reduction of Large-Scale Systems*, volume 45 of *Lecture Notes in Computational Science and Engineering*, pages 379–392. Springer, Heidelberg, 2005.

13. J. Daniel, W. B. Gragg, L. Kaufman, and G. W. Stewart. Reorthogonalization and stable algorithms for updating the Gram–Schmidt QR factorization. *Mathematics of Computation*, 30:772–795, 1976.

14. O. G. Ernst. Equivalent iterative methods for p-cyclic matrices. *Numer. Algorithms*, 25(1-4):161–180, 2000.

15. R. W. Freund, G. H. Golub, and N. M. Nachtigal. Recent advances in Lanczos-based iterative methods for nonsymmetric linear systems. In *Algorithmic trends in computational fluid dynamics (1991)*, ICASE/NASA LaRC Ser., pages 137–162. Springer, New York, 1993.

16. G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, MD, third edition, 1996.

17. R. Granat and B. Kågström. Direct eigenvalue reordering in a product of matrices in extended periodic Schur form. Report UMINF-05.05, Department of Computing Science, Umeå University, Umeå, Sweden, 2005.

18. S.J. Hammarling. Numerical solution of the stable, non-negative definite Lyapunov equation. *IMA J. Numer. Anal.*, 2:303–323, 1982.

19. J. J. Hench and A. J. Laub. Numerical solution of the discrete-time periodic Riccati equation. *IEEE Trans. Automat. Control*, 39(6):1197–1210, 1994.

20. I.M. Jaimoukha and E.M. Kasenally. Krylov subspace methods for solving large Lyapunov equations. *SIAM J. Numer. Anal.*, 31:227–251, 1994.

21. D. Kressner. The periodic QR algorithm is a disguised QR algorithm, 2003. To appear in *Linear Algebra Appl.*

22. D. Kressner. *Numerical Methods and Software for General and Structured Eigenvalue Problems*. PhD thesis, TU Berlin, Institut für Mathematik, Berlin, Germany, 2004.

23. R. B. Lehoucq. *Analysis and Implementation of an Implicitly Restarted Arnoldi Iteration*. PhD thesis, Rice U., 1995.

24. R. B. Lehoucq and D. C. Sorensen. Deflation techniques for an implicitly restarted Arnoldi iteration. *SIAM J. Matrix Anal. Appl.*, 17(4):789–821, 1996.

25. R. B. Lehoucq, D. C. Sorensen, and C. Yang. *ARPACK users' guide*. SIAM, Philadelphia, PA, 1998. Solution of large-scale eigenvalue problems with implicitly restarted Arnoldi methods.

26. W.-W. Lin, P. Van Dooren, and Q.-F. Xu. Periodic invariant subspaces in control. In *Proc. of IFAC Workshop on Periodic Control Systems, Como, Italy*, 2001.

27. K. Lust. *Numerical Bifurcation Analysis of Periodic Solutions of Partial Differential Equations*. PhD thesis, Department of Computer Science, KU Leuven, Belgium, 1997.

28. K. Lust. Improved numerical Floquet multipliers. *Internat. J. Bifur. Chaos Appl. Sci. Engrg.*, 11(9):2389–2410, 2001.

29. B. N. Parlett and J. Le. Forward instability of tridiagonal QR. *SIAM J. Matrix Anal. Appl.*, 14(1):279–316, 1993.

30. Y. Saad. *Numerical Methods for Large Eigenvalue Problems: Theory and Algorithms*. John Wiley, New York, 1992.

31. D. C. Sorensen. Implicit application of polynomial filters in a k-step Arnoldi method. *SIAM J. Matrix Anal. Appl.*, 13:357–385, 1992.

32. D. C. Sorensen. Implicitly restarted Arnoldi/Lanczos methods for large scale eigenvalue calculations. In *Parallel numerical algorithms (Hampton, VA, 1994)*, volume 4 of *ICASE/LaRC Interdiscip. Ser. Sci. Eng.*, pages 119–165. Kluwer Acad. Publ., Dordrecht, 1997.

33. J. Sreedhar and P. Van Dooren. Pole placement via the periodic Schur decomposition. In *Proceedings Amer. Contr. Conf.*, pages 1563–1567, 1993.

34. J. Sreedhar and P. Van Dooren. A Schur approach for solving some periodic matrix equations. In U. Helmke, R. Mennicken, and J. Saurer, editors, *Systems and Networks : Mathematical Theory and Applications*, volume 77, pages 339–362. Akademie Verlag, Berlin, 1994.

35. G. W. Stewart. *Matrix Algorithms. Vol. II.* SIAM, Philadelphia, PA, 2001. Eigensystems.

36. G. W. Stewart. A Krylov-Schur algorithm for large eigenproblems. *SIAM J. Matrix Anal. Appl.*, 23(3):601–614, 2001/02.

37. W. J. Stewart. *Introduction to the Numerical Solution of Markov Chains.* Princeton University Press, Princeton, NJ, 1994.

38. C. F. Van Loan. A general matrix eigenvalue algorithm. *SIAM J. Numer. Anal.*, 12(6):819–834, 1975.

39. A. Varga. Balancing related methods for minimal realization of periodic systems. *Systems Control Lett.*, 36(5):339–349, 1999.

40. A. Varga and P. Van Dooren. Computational methods for periodic systems - an overview. In *Proc. of IFAC Workshop on Periodic Control Systems, Como, Italy*, pages 171–176, 2001.

41. H. F. Walker. Implementation of the GMRES method using Householder transformations. *SIAM J. Sci. Stat. Comp.*, 9:152–163, 1988.

42. D. S. Watkins. Product eigenvalue problems. *SIAM Rev.*, 47:3–40, 2005.