

FAST HESSENBERG REDUCTION OF SOME RANK STRUCTURED MATRICES*

L. GEMIGNANI[†] AND L. ROBOL[‡]

Abstract. We develop two fast algorithms for Hessenberg reduction of a structured matrix $A = D + UV^H$, where D is a real or unitary $n \times n$ diagonal matrix and $U, V \in \mathbb{C}^{n \times k}$. The proposed algorithm for the real case exploits a two-stage approach by first reducing the matrix to a generalized Hessenberg form and then completing the reduction by annihilation of the unwanted subdiagonals. It is shown that the novel method requires $O(n^2k)$ arithmetic operations and is significantly faster than other reduction algorithms for rank structured matrices. The method is then extended to the unitary plus low rank case by using a block analogue of the CMV form of unitary matrices. It is shown that a block Lanczos-type procedure for the block tridiagonalization of $\Re(D)$ induces a structured reduction on A in a block staircase CMV-type shape. Then, we present a numerically stable method for performing this reduction using unitary transformations and show how to generalize the subdiagonal elimination to this shape, while still being able to provide a condensed representation for the reduced matrix. In this way the complexity still remains linear in k and, moreover, the resulting algorithm can be adapted to deal efficiently with block companion matrices.

Key words. Hessenberg reduction, quasi-separable matrices, bulge chasing, CMV matrix, complexity

AMS subject classification. 65F15

DOI. 10.1137/16M1107851

1. Introduction. Let $A = D + UV^H$ where D is a real or unitary $n \times n$ diagonal matrix and $U, V \in \mathbb{C}^{n \times k}$. Such matrices do arise commonly in the numerical treatment of structured (generalized) eigenvalue problems [1, 3]. We consider the problem of reducing A to upper Hessenberg form using unitary transformations, i.e., to find a unitary matrix Q such that $QAQ^H = H = (h_{ij})$ and $h_{ij} = 0$ for $i > j + 1$. Specialized algorithms exploiting the rank structure of A have complexity $O(n^2)$ whenever k is a small constant independent of n [12, 14]. However, for applications to generalized eigenproblems the value of k may be moderate or even large so that it is worthwhile to ask for the algorithm to be cost efficient w.r.t. the size k of the correction as well.

The Hessenberg reduction of a square matrix is the first basic step in computing its eigenvalues. Eigenvalue computation for (Hermitian) matrices modified by low rank perturbations is a classical problem arising in many applications [3]. More recently methods for diagonal plus low rank matrices have been used in combination with interpolation techniques in order to solve generalized nonlinear eigenvalue problems [1, 7]. Standard Hessenberg reduction algorithms for rank structured matrices [8, 12, 14] are both theoretically and practically ineffective as the size of the correction increases since their complexities depend quadratically or even cubically on k , or they suffer from possible instabilities. The aim of this paper is to describe a novel efficient reduction scheme which attains the cost of $O(n^2k)$ arithmetic operations (ops). For a

*Received by the editors December 13, 2016; accepted for publication (in revised form) by F. Tisseur May 2, 2017; published electronically June 22, 2017.

<http://www.siam.org/journals/simax/38-2/M110785.html>

Funding: The work of the authors was partially supported by the GNCS/INdAM project “Metodi numerici avanzati per equazioni e funzioni di matrici con struttura” and by the University of Pisa (grant PRA_2017_05).

[†]Dipartimento di Informatica, Università di Pisa, L.go Bruno Pontecorvo, 3, 56127 Pisa, Italy (l.gemignani@di.unipi.it).

[‡]ISTI, Area della ricerca CNR, Via G. Moruzzi 1, 56124 Pisa, Italy (leonardo.robol@isti.cnr.it).

real D the incorporation of this scheme into the fast QR-based eigensolver for rank-structured Hessenberg matrices given in [13] thus yields an eigenvalue method with overall cost $O(n^2k)$ ops. Further, in the unitary case the design of such a scheme is an essential step toward the development of an efficient eigenvalue solver for block companion matrices and, moreover, the techniques employed in the reduction are of independent interest for computations with matrix valued orthogonal polynomials on the unit circle [4, 22, 23].

The proposed algorithm consists of two stages including a first intermediate reduction to a generalized τ_k -Hessenberg form, where $\tau_k = k$ in the real case and $\tau_k = 2k$ in the unitary case. A two-stage approach has been shown to be beneficial for blocked and parallel computations [9, 20]. Our contribution is to show that the same approach can conveniently be exploited in the framework of rank-structured computations. Specifically, to devise a fast reduction algorithm we find Q as the product of two unitary matrices $Q = Q_2Q_1$ determined as follows:

Reduction to banded form. $Q_1DQ_1^H$ is a banded matrix of bandwidth τ_k , and Q_1U is upper triangular. In particular, this implies that the matrix $B := Q_1AQ_1^H$ is in generalized τ_k -Hessenberg form, that is, its elements are 0 below the τ_k th subdiagonal.

Subdiagonal elimination. We compute a unitary matrix Q_2 such that $Q_2BQ_2^H$ is in upper Hessenberg form. The process employs a sequence of Givens rotations used to annihilate the last $\tau_k - 1$ subdiagonals of the matrix B .

It is shown that both steps can be accomplished in $O(n^2k)$ ops. For a real D the algorithm for the reduction to banded form relies upon an adaptation of the scheme proposed in [3] for bordered matrices. The subdiagonal annihilation procedure employed at the second step easily takes advantage of the data-sparse representation of the matrices involved in the updating process. The extension to the unitary setting is much more cumbersome and needs some additional results. By exploiting the relationships between the computation at stage 1 and the block Lanczos process we prove that the band reduction scheme applied to $\Re(D)$ and $U_D := [U, DU]$ computes a unitary matrix Q_1 such that $F := Q_1DQ_1^H$ is $2k$ -banded in staircase form. A suitable reblocking of F in a block tridiagonal form puts in evidence analogies and generalizations of this form with the classical CMV format of unitary matrices [10, 11, 21], and this is why we refer to F as a block CMV-type matrix. An application of CMV matrices for polynomial rootfinding is given in [5]. The block CMV structure plays a fundamental role for reducing the overall cost of the reduction process. To this aim we first introduce a numerically stable algorithm for computing F in factored form using a suitable generalization of the classical Schur parametrization of scalar CMV matrices [10]. Then, we explain how bulge-chasing techniques can be efficiently employed in the subsequent subdiagonal elimination stage applied to the cumulative B to preserve the CMV structure of the part of the matrix that actually needs to be still reduced. Some experimental results are finally presented to illustrate the speed and the accuracy of our algorithms.

The paper is organized as follows. The real case is treated in section 2. The two reduction steps applied to a real input matrix D are discussed and analyzed in subsections 2.1 and 2.2, respectively. The generalization for unitary matrices is presented in section 3. In particular, subsections 3.3 and 3.5 deal with the construction of the block CMV structure in the reduction process, whereas the computation and the condensed representation of the final Hessenberg matrix are the topics of subsections 3.6 and 3.7. Experimental results are described in section 4 followed by some conclusions and future work in section 5.

2. Reduction processes: The real plus low rank case. In this section we present a two-phase Hessenberg reduction algorithm for a matrix $A = D + UV^H$, where D is a real $n \times n$ diagonal matrix and $U, V \in \mathbb{C}^{n \times k}$.

2.1. Reduction to banded form. First, we show how it is possible to transform the matrix A into a Hermitian banded matrix plus a low rank correction, i.e., $Q_1 A Q_1^H = D_1 + U_1 V_1^H$, where

$$D_1 = \begin{bmatrix} \times & \dots & \times & & \\ \vdots & \ddots & & \ddots & \\ \times & & \ddots & & \times \\ & \ddots & & \ddots & \vdots \\ & & \times & \dots & \times \end{bmatrix}, \quad U_1 = \begin{bmatrix} X \\ 0 \end{bmatrix}, \quad X \in \mathbb{C}^{k \times k},$$

and D_1 has bandwidth k . The computation can equivalently be reformulated as a band reduction problem for the bordered matrix

$$E = \left[\begin{array}{c|c} O_k & V^H \\ \hline U & D \end{array} \right],$$

where O_k denotes the zero matrix of size k . The algorithm we present here is basically an adaptation of [2, Algorithm 2] for solving this problem efficiently. It finds a unitary matrix $Q = I_k \oplus Q_1$ such that QEQ^H is in generalized k -Hessenberg form.

The main idea behind the reduction is that the matrix D at the start is a (very special kind of) k -banded matrix, since it is diagonal. We then start to put zeros in the matrix U using Givens rotations, and we right-multiply by the same rotations to obtain a unitary transformation. This, generally, will degrade the banded structure of the matrix D . We then restore it by means of Givens rotations acting on the left in correspondence of the zeros of U . Since the rotations must leave the zeros unchanged, this will also preserve the partial structure of U .

More precisely, let $\ell(i, j)$ be the function defined as

$$\begin{aligned} \ell(i, j) : \quad \mathbb{N}^2 &\longrightarrow \mathbb{N}, \\ (i, j) &\longmapsto k(n + j - i) + j. \end{aligned}$$

LEMMA 1. *The function ℓ is injective from the subset $I \subseteq \mathbb{N}^2$ defined by $I = [1, n] \times [1, k] \cap \mathbb{N}^2$.*

Proof. Assume that $\ell(i, j) = \ell(i', j')$. Then we have that

$$0 = \ell(i, j) - \ell(i', j') = k((j - j') - (i - i')) + j - j'.$$

This implies that $j \equiv j'$ modulo k , but since $j \in [1, k]$ we can conclude that $j = j'$. Back substituting this in the above equation yields that $i = i'$, thus giving the thesis. \square

Since ℓ is injective and \mathbb{N} is totally ordered the map ℓ induces an order on the set I . This is exactly the order in which we will put the zeros in the matrix U . Pictorially,

this can be described by the following:

$$\begin{bmatrix} \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \end{bmatrix} \rightarrow \begin{bmatrix} \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \\ 0 & \times & \times \end{bmatrix} \rightarrow \begin{bmatrix} \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \\ 0 & \times & \times \\ 0 & \times & \times \end{bmatrix} \rightarrow \begin{bmatrix} \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \\ 0 & \times & \times \\ 0 & 0 & \times \end{bmatrix} \rightarrow \begin{bmatrix} \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \\ \times & \times & \times \\ 0 & \times & \times \\ 0 & \times & \times \\ 0 & 0 & \times \end{bmatrix} \rightarrow \dots$$

The annihilation scheme proceeds by introducing the zeros from the bottom up along downwardly sloping diagonals. Let $\mathcal{R}(A, (k, j)) = G_k A G_k^H$, where $G_k = I_{k-2} \oplus \mathcal{G}_k \oplus I_{n-k}$ is a Givens rotation in the $(k-1, k)$ -plane that annihilates $a_{k,j}$. The band reduction algorithm is then described as follows:

```

Algorithm Band_Red
for  $i = 1 : n - 1$ 
  for  $j = 1 : i$ 
     $A = \mathcal{R}(A, \ell^{-1}(ki + j));$ 
  end
end
    
```

It is easily seen that we can perform each similarity transformation in $O(k)$ arithmetic operations and, therefore, the overall cost of the band reduction is $O(n^2k)$.

2.2. Subdiagonal elimination. Our algorithm for subdiagonal elimination and Hessenberg reduction exploits the properties of bulge-chasing techniques applied to a band matrix. Observe that at the very beginning of this stage the input matrix A can be reblocked as a block tridiagonal matrix plus a rank- k correction located in the first block row. At the first step of the Hessenberg reduction we determine a unitary matrix $\mathcal{H}_1 \in \mathbb{C}^{k \times k}$ such that $H_1 A(2 : k + 1, 1)$ is a multiple of the first unit vector. Then the transformed matrix $A := H_1 A H_1^H$, where $H_1 = 1 \oplus \mathcal{H}_1 \oplus I_{n-k-1}$, reveals a bulge outside the current band which has to be chased off before proceeding with the annihilation process. In the next figure we illustrate the case where $k = 3$.

$$\begin{array}{cccccccc} \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times \\ \times & \times & \times & \times & \times & \times & \times & \times \\ 0 & \times & \times & \times & \times & \times & \times & \times \\ 0 & 0 & \times & \times & \times & \times & \times & \times \\ 0 & 0 & 0 & \times & \times & \times & \times & \times \\ 0 & 0 & 0 & 0 & \times & \times & \times & \times \end{array} \rightarrow \begin{array}{cccccccc} \times & \times & \times & \times & \times & \times & \times & \times \\ 0 & \times & \times & \times & \times & \times & \times & \times \\ 0 & \times & \times & \times & \times & \times & \times & \times \\ 0 & \times & \times & \times & \times & \times & \times & \times \\ 0 & \times & \times & \times & \times & \times & \times & \times \\ 0 & \times & \times & \times & \times & \times & \times & \times \\ 0 & \times & \times & \times & \times & \times & \times & \times \\ 0 & 0 & 0 & 0 & \times & \times & \times & \times \end{array}$$

The crucial observation is that the bulge inherits the rank structure of the matrix \mathcal{H}_1^H . If we construct \mathcal{H}_1 as product of consecutive Givens rotations, we find that \mathcal{H}_1^H is unitary lower Hessenberg and therefore the lower triangular part of the bulge has rank one at most. Then the bulge can be moved down by computing its QR decomposition. From the rank property it follows that the Q factor is again a unitary Hessenberg matrix so that the process can be continued until the bulge disappears. By swapping the elementary Givens rotations employed in both the annihilation and bulge-chasing step we arrive at the following algorithm.

```

Algorithm Subd_El
for  $j = 1 : n - 2$ 
  for  $i = \min\{n, j + k\} : -1 : j + 2$ 
     $A = \mathcal{R}(A, (i, j));$ 
  for  $s = i + k : k : n$ 
     $A = \mathcal{R}(A, (s, s - k - 1));$ 
  end
end
end
    
```

The arithmetic cost of the algorithm is upper bounded by $\lceil \frac{n}{k} \rceil nkp$, where p is the maximum number of arithmetic operations required to perform the updating $A \rightarrow A := \mathcal{R}(A, (i, j))$ at each step of the algorithm. A data-sparse representation of the matrix A is provided by the following.

LEMMA 2. *At each step of the algorithm the matrix A satisfies*

$$A - A^H = UV^H - VU^H$$

for suitable $U, V \in \mathbb{C}^{n \times k}$.

Proof. The input matrix $A = D_1 + U_1V_1^H$ satisfies $A - A^H = U_1V_1^H - V_1U_1^H$, and the relation is maintained under the congruence transformation $A \rightarrow A := \mathcal{R}(A, (i, j))$. \square

Let $\Phi: \mathbb{C}^{n \times n} \times \mathbb{C}^{n \times k} \times \mathbb{C}^{n \times k} \rightarrow \mathbb{C}^{n \times n}$ be defined as follows:

$$\Phi(A, U, V) = \text{tril}(A) + (\text{tril}(A, -1))^H + (\text{triu}(UV^H - VU^H, 1)).$$

Based on the previous result we know that the above map provides an efficient parametrization of the matrix A , so we can replace the (unstructured) transformation $A \rightarrow A := \mathcal{R}(A, (i, j))$ in Algorithm Algorithm Subd_El with its condensed variant

$$\begin{aligned}
 (\text{tril}(A), U, V) &\rightarrow (\text{tril}(A), U, V) := \mathcal{R}_\Phi(\Phi(\text{tril}(A), U, V), (i, j)), \\
 \Phi(\mathcal{R}_\Phi(\Phi(\text{tril}(A), U, V), (i, j))) &= \mathcal{R}(\Phi(\text{tril}(A), U, V), (i, j)) = \mathcal{R}(A, (i, j)),
 \end{aligned}$$

which can be employed by updating only the generators of the data-sparse representation at the cost of $p = O(k)$ operations, thus yielding the overall cost of $O(n^2k)$ ops.

3. Reduction processes: The unitary plus low rank case. In this section we deal with an extension of the approach to the case of unitary plus low rank matrices. More precisely, we consider the case where $A = D + UV^H$ with D being an $n \times n$ unitary diagonal matrix and $U, V \in \mathbb{C}^{n \times k}$.

A strong motivation for the interest in this class of matrices comes from the companion block forms. If $A(x) = \sum_{i=0}^n A_i x^i$ with $A_i \in \mathbb{C}^{m \times m}$ is a matrix polynomial, then the linear pencil

$$L(x) = x \begin{bmatrix} A_n & & & \\ & I_m & & \\ & & \ddots & \\ & & & I_m \end{bmatrix} - \begin{bmatrix} -A_{n-1} & \dots & -A_1 & -A_0 \\ & I_m & & \\ & & \ddots & \\ & & & I_m \end{bmatrix}$$

has the same spectral structure of $A(x)$ [18]. When the polynomial is monic the same holds for the above pencil (and so the generalized eigenvalue problem associated with it turns into a standard eigenvalue problem) and the constant term can be written as $Z + UW^H$, where $U = e_1 \otimes I_m$, $W^H = [-A_{n-1}, \dots, -A_0 - I]$ and $Z = Z_n \otimes I_m$, where Z_n is the $n \times n$ shift matrix that maps e_n to e_1 and e_i to e_{i+1} for $i < n$. In particular, Z_n can be easily diagonalized so that we can consider an equivalent eigenvalue problem for the matrix $D + \tilde{U}\tilde{W}^H$, where D is unitary diagonal generated from the n th roots of unity.

3.1. Properties of unitary plus low rank matrices. The results stated in the previous section mainly relied upon the fact that the matrix D is real in order to guarantee that, at each step of the reduction, the transformed matrix PDP^H is Hermitian, where P denotes the product of all Givens rotations accumulated so far. In particular, for a real D maintaining the matrix PDP^H lower banded ensures that the whole matrix is also upper banded (thanks to the Hermitian structure). This is not true anymore in the unitary case. In fact, applying the previous algorithm “as is” leads to a matrix PDP^H which is in generalized k -Hessenberg form, that is, it is lower banded (with bandwidth k) but generally dense in its upper triangular part.

It is possible to prove, relying on the formalism of quasi-separable matrices, that a structure is indeed present in the matrix PDP^H , even if it is not readily available to be exploited in numerical computations. Let us first recall some preliminaries. (See [15, 16] for a survey of properties of quasi-separable matrices.)

DEFINITION 3. Let $A \in \mathbb{C}^{n \times n}$ a matrix. We say that A is (k_l, k_u) -quasi-separable if, for any $i = 1, \dots, n - 1$, we have

$$\text{rank}(A[1 : i, i + 1 : n]) \leq k_u, \quad \text{rank}(A[i + 1 : n, 1 : i]) \leq k_l,$$

where we have used the MATLAB notation to identify off-diagonal submatrices of A .

Banded matrices with lower bandwidth k_l and upper bandwidth k_u are a subset of (k_l, k_u) -quasi-separable matrix. This is a direct consequence of the fact that, in a banded matrix, each off-diagonal submatrix has all the nonzero entries in the top-right (or bottom-left) corner, and its rank is bounded by the bandwidth of the large matrix. However, quasi-separable matrices enjoy some nice properties that banded matrices do not have, such as closure under inversion.

Before going into the details we introduce a very simple but important result that is relevant to this setting, called the nullity theorem [17, 19].

THEOREM 4 (nullity). Suppose $A \in \mathbb{C}^{n \times n}$ is a nonsingular matrix and let α and β to be nonempty proper subsets of $\mathbb{I}_n := \{1, \dots, n\}$. Then

$$\text{rank}(A^{-1}[\alpha; \beta]) = \text{rank}(A[\mathbb{I}_n \setminus \beta; \mathbb{I}_n \setminus \alpha]) + |\alpha| + |\beta| - n,$$

where, as usual, $|J|$ denotes the cardinality of the set J . In particular, if A and A^{-1} are partitioned conformally as

$$A = \begin{bmatrix} A_{1,1} & A_{1,2} \\ A_{2,1} & A_{2,2} \end{bmatrix}, \quad A^{-1} = \begin{bmatrix} B_{1,1} & B_{1,2} \\ B_{2,1} & B_{2,2} \end{bmatrix}$$

with square diagonal blocks, then the rank of each off-diagonal block of A matches the rank of the corresponding block in the inverse.

Hermitian and unitary quasi-separable matrices share an additional property, which we call *rank symmetry*, which makes them suitable to develop fast algorithms.

THEOREM 5. *Let U be a unitary or Hermitian $n \times n$ matrix. Then if U is lower quasi-separable of rank k it is also upper quasi-separable with the same rank.*

Proof. If U is Hermitian the result is obvious since each off-diagonal submatrix in the lower triangular part corresponds to another submatrix in the upper triangular part. In the unitary case, instead, we can rely on the nullity theorem that guarantees that each submatrix of the form $P[i+1 : n, 1 : i]$ has rank equal to the same submatrix in the inverse of P . Since $P^{-1} = P^H$ we have the thesis. \square

Remark 6. Let $A = D + UV^H$, where D is unitary and $U, V \in \mathbb{C}^{n \times k}$, and consider a k -Hessenberg¹ form of A defined by

$$H = PAP^H = PDP^H + (PU)(PV)^H.$$

Suppose that PU is upper triangular. Since PDP^H is rank symmetric and $H - (PU)(PV)^H$ has lower bandwidth k we can conclude that the upper quasi-separable rank of H is bounded by $\tau_k = 2k$.

3.2. The block CMV structure. The results of the previous section are not constructive nor computationally helpful (at least in this form), since they do not provide a representation for the upper part of the matrix PDP^H and a fortiori of the cumulative matrix H . In the spirit of [6] we pursue a different approach by showing how to induce a particular banded structure of width $2k$ (the so-called CMV shape [5, 21]) on the unitary matrix PDP^H . In this form the quasi-separable structure of PDP^H , even if not of minimal size, will be evident by the banded structure and, therefore, easily recoverable and representable.

DEFINITION 7 (CMV shape). *A unitary matrix A is said to be CMV structured with block size k if there exist $k \times k$ nonsingular matrices R_i and L_i , respectively, upper and lower triangular, such that*

$$A = \begin{bmatrix} \times & \times & L_3 & & & \\ R_1 & \times & \times & & & \\ & \times & \times & \times & L_5 & \\ & R_2 & \times & \times & \times & \\ & & & \times & \times & \\ & & & R_4 & \times & \ddots \\ & & & & & \ddots & \ddots \end{bmatrix},$$

where the symbol \times has been used to identify (possibly) nonzero blocks.

In order to simplify the notation we often assume that n is a multiple of $2k$, so the above structures fit “exactly” in the matrix. However, this is not restrictive and the theory presented here continues to hold in greater generality. In practice, one can deal with the more general case by allowing the blocks in the bottom-right corner of the matrix to be smaller.

Remark 8. Notice that a matrix in CMV form with blocks of size k is, in particular, $2k$ -banded. The CMV structure with blocks of size 1 has been proposed as a generalization of what the tridiagonal structure is for Hermitian matrices in [11] and [21]. A further analogy between the scalar and the block case is derived from the

¹By k -Hessenberg we mean a matrix with at most k subdiagonals different from zero.

nullity theorem. We have for $p > 0$

$$0 = \text{rank}(A[1 : 2pk, 2(p+1)k + 1 : n]) = \text{rank}(A[2pk + 1 : n, 1 : 2(p+1)k]) - 2k,$$

which gives

$$\text{rank}(A[2pk + 1 : 2(p+1)k, (2p-1)k + 1 : (2p+1)k]) = k.$$

Pictorially we are setting rank constraints on the blocks

$$A = \begin{bmatrix} \times & \times & L_3 & & & \\ R_1 & \times & \times & & & \\ & \boxed{\times} & \boxed{\times} & \times & L_5 & \\ & \boxed{\times} & \boxed{\times} & \times & \times & \\ & & & \boxed{\times} & \boxed{\times} & \ddots \\ & & & \boxed{\times} & \boxed{\times} & \ddots \\ & & & & & \ddots \end{bmatrix}$$

and by similar arguments on the corresponding blocks in the upper triangular portion.

3.3. Existence of the CMV structure. This section is devoted to showing that the reduction process simultaneously applied to the diagonal matrix $\mathfrak{R}(D)$ and the rectangular one $U_D = [U, DU]$ induces a CMV structure on the matrix $A = D + UV^H$. Observe that the construction of a unitary P such that $T_D := P\mathfrak{R}(D)P^H$ is $2k$ -banded and, moreover, PU_D is upper triangular can theoretically be analyzed in the framework of block Krylov (Lanczos) methods for converting $\mathfrak{R}(D)$ in block tridiagonal form having fixed the first block column of P^H . Exact breakdowns occur when the iteration cannot be continued due to a rank drop and some modifications are required to generate a block tridiagonal T_D with subdiagonal blocks of varying (decreasing) sizes. Conversely, we characterize the regular case where there is no breakdown (and no rank reduction) in the corresponding block Lanczos process by stating the following.

DEFINITION 9 (no breakdown condition). *We say that the matrix $A \in \mathbb{C}^{n \times n}$ satisfies the no-breakdown-condition related to $U \in \mathbb{C}^{n \times p}$ and $P \in \mathbb{C}^{n \times n}$ unitary if PU is upper triangular of full rank and $H = PAP^H = (H_{i,j})$ is a block upper Hessenberg matrix, $H_{i,j} \in \mathbb{C}^{p \times p}$, $1 \leq i, j \leq s$, and $H_{j+1,j}$ with $1 \leq j \leq s-1$ are of full rank.*

In order to simplify our analysis of the reduction process simultaneously applied to the diagonal matrix $\mathfrak{R}(D)$ and the rectangular one $U_D = [U, DU]$ we assume the following two conditions are satisfied:

1. The matrix $\mathfrak{R}(D)$ satisfies the no-breakdown-condition related to U_D and the unitary matrix $Q^H \in \mathbb{C}^{n \times n}$ computed at end of the band reduction process such that $Q^H \mathfrak{R}(D)Q = T_D$ is block tridiagonal.
2. The dimension n of $\mathfrak{R}(D)$ is a multiple of $2k$, i.e., $n = 2k\ell$. This requirement ensures that all the blocks of T_D have constant size $p = 2k$.

We notice that the requirement (1) is not so strict as it might seem. In fact, whenever a breakdown happens or, equivalently speaking, we find a rank loss in the associated Lanczos scheme, then the reduction process can be continued in such a way that the matrix obtained at the very end will still have a block CMV structure with blocks

of possibly varying sizes. However, this would complicate the handling of indices to track this situation, and so we will not cover this case explicitly.

The next results provides our algorithm for the reduction of a unitary diagonal D into a block CMV shape.

LEMMA 10. *Let $B_R = Q^H \mathfrak{R}(D)Q$ be the $2k$ -banded matrix obtained by applying the reduction scheme to the diagonal matrix $\mathfrak{R}(D)$ and the rectangular one $[U \ DU]$. Assume also that the no-breakdown-condition of $\mathfrak{R}(D)$ related to $[U \ DU]$ and Q^H is satisfied. Then the following structures are induced on B_I and B defined as follows:*

- (i) $B_I = Q^H \mathfrak{S}(D)Q$ is $2k$ -banded.
- (ii) $B = Q^H DQ$ is $2k$ -banded and CMV structured with blocks of size k .

Proof. We can look at the $2k$ -banded structure as a characterization of the operator $\mathfrak{R}(D)$ acting on \mathbb{C}^n . In particular, it tells us that, if we see the columns of Q as a basis for \mathbb{C}^n , the j th one is mapped by $\mathfrak{R}(D)$ in the subspace spanned by the first $j + 2k$ ones. Since, by construction, we have that

$$Q^H [U \ DU] = R \iff [U \ DU] = Q \begin{bmatrix} \tilde{R} \\ 0_{2k} \\ \vdots \\ 0_{2k} \end{bmatrix}, \quad \tilde{R} \text{ nonsingular,}$$

we have that the first $2k$ columns of Q are a basis for the space generated by the columns of $[U \ DU]$. Under the no breakdown condition we have that the $(2k + 1)$ th column of Q will be a multiple of the part of $\mathfrak{R}(D)Ue_1$ orthogonal to $[U \ DU]$. Notice that $\mathfrak{R}(D) = \frac{1}{2}(D + D^H)$ and therefore the $(2k + 1)$ th column of Q will be also a multiple of the part of $D^H Ue_1$ orthogonal to $[U \ DU]$. Extending this remark yields that Q is the unitary factor in a QR factorization of a Krylov-type matrix \tilde{U} , that is,

$$(1) \quad QR_U = \underbrace{\begin{bmatrix} U & DU & D^H U & D^2 U & (D^2)^H U & D^3 U & \dots & (D^{\ell-1})^H U & D^\ell U \end{bmatrix}}_{\tilde{U}}$$

with R_U invertible. Since $\mathfrak{S}(D) = \frac{1}{2i}(D - D^H)$ we have that $Q^H \mathfrak{S}(D)Q$ has a $2k$ -banded structure and the j th column of Q is mapped by the linear operator $\mathfrak{S}(D)$ into the span of the first $j + 2k$ ones. Given that $D = \mathfrak{R}(D) + i\mathfrak{S}(D)$ we have that also $Q^H DQ$ is $2k$ -banded.

It remains to show that $Q^H DQ$ is in fact endowed with a block CMV structure. For this it is convenient to rely upon (1). We can characterize the i th block column of B by using

$$BE_i = Q^H DQE_i \Rightarrow QBE_i = D\tilde{U}R_U^{-1}E_i,$$

where we set $E_i := e_i \otimes I_k$ and $D\tilde{U}$ can be explicitly written as

$$D\tilde{U} = \begin{bmatrix} DU & D^2 U & U & D^3 U & D^H U & D^4 U & \dots & (D^{\ell-2})^H U & D^{\ell+1} U \end{bmatrix}.$$

Recall that R_U^{-1} is upper triangular, and each block column of $\tilde{U}R_U^{-1}$ is a combination of the corresponding one in \tilde{U} and the previous ones. Thus we can deduce the following:

- The block column of $\tilde{U}R_U^{-1}$ corresponding to the ones of the form $D^j U$ in \tilde{U} (which are in position $2j$, except for the first one) are mapped into a linear combination of the first $2(j + 1)$ block columns by the action of D on the left.

- Since the block column of the form $(D^H)^j U$ of \tilde{U} are mapped in $(D^H)^{j-1} U$ the corresponding block column in $\tilde{U} R_U^{-1}$ will be mapped in a combination of the previous ones.

The above remarks give bounds on the maximum lower band structure that are stricter than the ones given by being $2k$ -banded, and have the following shape:

$$Q^H D Q = \begin{bmatrix} \times & \times & \times & \times & \times \\ R_1 & \times & \times & \times & \times \\ & \times & \times & \times & \times \\ & & R_2 & \times & \times \\ & & & \times & \times \\ & & & & R_4 & \times \end{bmatrix},$$

where R_i are suitable nonsingular upper triangular matrices. Performing a similar analysis on $(Q^H D Q)^H = Q^H D^H Q$ yields a similar structure also in the upper part, proving our claim about the CMV block structure of $Q^H D Q$. \square

The above description is satisfying theoretically, but applying the Lanczos process to implicitly obtain the block CMV structure is not ideal numerically. In the next subsection we introduce a different factored form of block CMV structured matrices which is the basis of a numerically stable algorithm for CMV reduction presented in subsection 3.5.

3.4. Factorizing CMV matrices. Block CMV matrices have interesting properties that we will exploit in the following to perform the reduction. We show here that they can be factorized in block diagonal matrices with $2k \times 2k$ diagonal blocks.

THEOREM 11. *Let $A \in \mathbb{C}^{n \times n}$, $n = 2k\ell$, be a block CMV matrix as in Definition 7. Then there exist two block diagonal unitary matrices A_1 and A_2 such that $A = A_1 A_2$, which have the following form:*

$$A_1 = \begin{bmatrix} \times & L_{A,2} & & & \\ R_{A,1} & \times & & & \\ & & \ddots & & \\ & & & \times & L_{A,2\ell} \\ & & & R_{A,2\ell-1} & \times \end{bmatrix}, \quad A_2 = \begin{bmatrix} I_k & & & & \\ & \times & L_{A,3} & & \\ & R_{A,2} & \times & & \\ & & & \ddots & \\ & & & & \times \end{bmatrix},$$

where the matrices $R_{A,j}$ are upper triangular and $L_{A,j}$ are lower triangular. Moreover, for any A_1 and A_2 with the above structure, the matrix A is block CMV.

Proof. The proof that we provide is constructive and gives an algorithm for the computation of A_1 and A_2 . Assume that A has the structure

$$A = \begin{bmatrix} A_{1,1} & A_{1,2} & L_3 & & \\ R_1 & A_{2,2} & A_{2,3} & & \\ & \vdots & \vdots & \ddots & \end{bmatrix},$$

and let $Q_1 \begin{bmatrix} S_1 \\ 0 \end{bmatrix} = \begin{bmatrix} A_{1,1} \\ R_1 \end{bmatrix}$ be a QR decomposition of the first block column. Then we can factor A as

$$\begin{bmatrix} Q_1^H \\ I \end{bmatrix} A = \begin{bmatrix} S_1 & \tilde{A}_{1,2} & \tilde{A}_{1,3} \\ & \tilde{A}_{2,2} & \tilde{A}_{2,3} \end{bmatrix}.$$

Since S_1 is upper triangular and A is unitary we must have that S_1 is diagonal and $\tilde{A}_{1,2} = \tilde{A}_{1,3} = 0$ because of the nullity theorem. Thus, we can assume to choose Q_1 so that $S_1 = I$ and we obtain

$$A = \begin{bmatrix} Q_1 & & & \\ & I & & \\ & & \ddots & \\ & & & I \end{bmatrix} \begin{bmatrix} I & & & & \\ \tilde{A}_{2,2} & \tilde{A}_{2,3} & & & \\ A_{3,2} & A_{3,3} & A_{3,4} & R_5^H & \\ R_2 & A_{4,3} & A_{4,4} & A_{4,5} & \\ & & \vdots & & \ddots \end{bmatrix}.$$

Notice that, if we look at Q_1 as a 2×2 block matrix, the block in position $(2,1)$ has to be upper triangular, and we can force the one in position $(1,2)$ to be lower triangular. (The last k columns are only determined up to right multiplication by a $k \times k$ invertible matrix.)

We can continue the procedure by computing a QR factorization

$$Q_2 \begin{bmatrix} S_2 \\ 0 \end{bmatrix} = \begin{bmatrix} A_{3,2} \\ R_2 \end{bmatrix} \implies Q_2^H \begin{bmatrix} A_{3,2} & A_{3,3} \\ R_2 & A_{4,3} \end{bmatrix} = \begin{bmatrix} S_2 & \times \\ 0 & 0 \end{bmatrix}$$

since the right-hand side has rank k in view of Remark 8 and we assume R_2 to be of full rank. This provides a new factorization

$$A = \begin{bmatrix} Q_1 & & & \\ & Q_2 & & \\ & & \ddots & \\ & & & I \end{bmatrix} \begin{bmatrix} I & & & & \\ \tilde{A}_{2,2} & \tilde{A}_{2,3} & & & \\ S_2 & \times & & & \\ & & \tilde{A}_{4,4} & \tilde{A}_{4,5} & \\ & & \vdots & & \ddots \end{bmatrix},$$

where the upper zero structure follows from the lower one by reapplying the nullity theorem again and from the fact that the triangular matrices are nonsingular. Iterating this procedure until the end provides a factorization of the form:

$$A = \begin{bmatrix} \times & L_{A,2} & & & \\ R_{A,1} & \times & & & \\ & & \ddots & & \\ & & & \times & L_{A,2\ell} \\ & & & R_{A,2\ell-1} & \times \end{bmatrix} \cdot \begin{bmatrix} I_k & & & & \\ & \times & \times & & \\ & R_{A,2} & \times & & \\ & & & \ddots & \\ & & & & \times \end{bmatrix}.$$

As a last step, we observe that the superdiagonal blocks in the right factor need to be lower triangular because of the presence of the lower triangular blocks in A and for the full-rank condition that we have imposed on the triangular block entries. Therefore, the previous factorization has the required form. The other implication can be checked by a direct computation. □

3.5. From diagonal to CMV. We are now concerned with the transformation of a diagonal plus low-rank matrix $A = D + UV^H$, with D unitary, to a CMV plus low-rank one where the block vector U is upper triangular. We have proved in section 3.3 that this transformation can virtually be performed under the no-breakdown condition. Here we are going to present a different numerically stable approach which works without restrictions and outputs a factored representation of

the block CMV matrix with relaxed constraints on the invertibility of the triangular blocks.

In order to perform this reduction, we shall introduce some formal concepts that will greatly ease the handling of the structure.

DEFINITION 12. We say that a unitary matrix $Q_j \in \mathbb{C}^{n \times n}$ is a block unitary transformation with block size k if there exist an integer s and a $2k \times 2k$ unitary matrix \hat{Q}_j such that

$$Q_j = \begin{bmatrix} I_{(j-1)k} & & \\ & \hat{Q}_j & \\ & & I_{sk} \end{bmatrix}.$$

The matrix \hat{Q}_j is called the active block of Q_j , and the integer $j = \frac{n}{k} - s - 1$ is used to denote its position on the block diagonal.

Informally, we can think of Q_j as a block version of Givens rotations or, more generally, of essentially 2×2 unitary matrices.

To illustrate our interest in block unitary transformations we give the following simple application, which will be the base of our reduction process.

LEMMA 13. Let $U \in \mathbb{C}^{n \times k}$, $n = \ell k$. Then there exists a sequence of block unitary transformations $Q_1, \dots, Q_{\ell-1}$ such that

$$Q_1 \dots Q_{\ell-1} U = R$$

with R being an upper triangular matrix.

Block unitary transformations have some useful properties that we will exploit to make our algorithm asymptotically fast.

LEMMA 14. Let $A_j, B_j, C_j \in \mathbb{C}^{n \times n}$ be block unitary transformations of block size k and D a unitary diagonal matrix of size $n = \ell k$ for some ℓ . Then, the following properties hold:

- (i) If $|j - i| > 1$, then $A_j A_i = A_i A_j$.
- (ii) There exists a block unitary transformation E_i such that $E_i = A_i B_i$.
- (iii) Given the sequence $A_1, \dots, A_{\ell-1}$, there exists a modified sequence of transformations $\tilde{A}_1, \dots, \tilde{A}_{\ell-1}$ such that $\tilde{A}_1 \dots \tilde{A}_{\ell-1} = A_1 \dots A_{\ell-1} D$.
- (iv) For any choice of A_j, B_{j+1}, C_j with “V-shaped” indices, there exist three block unitary transformations $\tilde{A}_{j+1}, \tilde{B}_j, \tilde{C}_{j+1}$ such that $A_j B_{j+1} C_j = \tilde{A}_{j+1} \tilde{B}_j \tilde{C}_{j+1}$.

Proof. All the properties can be checked easily using the definition of block unitary transformation. We only discuss (iv). Notice that the product of the block unitary transformations is a $3k \times 3k$ unitary matrix embedded in an identity. Assume for simplicity that the size of the matrices is $3k$, so we do not have to keep track of the identities.

Let $S := A_1 B_2 C_1$. We can compute its QR factorization by using 3 block unitary transformations such that

$$\tilde{A}_2^H S = \begin{bmatrix} \times & \times & \times \\ \times & \times & \times \\ 0_k & \times & \times \end{bmatrix}, \quad \tilde{B}_1^H \tilde{A}_2^H S = \begin{bmatrix} \times & \times & \times \\ 0_k & \times & \times \\ 0_k & \times & \times \end{bmatrix}, \quad \tilde{C}_2^H \tilde{B}_1^H \tilde{A}_2^H S = D,$$

where D is a diagonal unitary matrix (since an upper triangular unitary matrix has to be diagonal). The diagonal term can be absorbed into the block unitary transformations, so we can assume without loss of generality that $D = I$. Thus, we have the new decomposition $A_1 B_2 C_1 = S = \tilde{A}_2 \tilde{B}_1 \tilde{C}_2$, as desired. \square

A strict relationship exists between block CMV matrices and block unitary transformations.

LEMMA 15. *A unitary matrix $A \in \mathbb{C}^{n \times n}$, $n = \ell k$, is CMV structured with block size k if and only if there exist block unitary transformations $\mathcal{A}_1, \dots, \mathcal{A}_{\ell-1}$ such that*

$$A = \begin{cases} \mathcal{A}_1 \mathcal{A}_3 \dots \mathcal{A}_{\ell-1} \mathcal{A}_2 \mathcal{A}_4 \dots \mathcal{A}_{\ell-2} & \text{if } \ell \text{ is even,} \\ \mathcal{A}_1 \mathcal{A}_3 \dots \mathcal{A}_{\ell-2} \mathcal{A}_2 \mathcal{A}_4 \dots \mathcal{A}_{\ell-1} & \text{if } \ell \text{ is odd,} \end{cases}$$

and all the block unitary transformations are of the form

$$\mathcal{A}_j = \begin{bmatrix} \times & L \\ R & \times \end{bmatrix}, \quad R, L^H \text{ nonsingular upper triangular.}$$

Proof. Assume that ℓ is even. The other case is handled in the same way. Let $A_1 := \mathcal{A}_1 \mathcal{A}_3 \dots \mathcal{A}_{\ell-1}$ and $A_2 := \mathcal{A}_2 \mathcal{A}_4 \dots \mathcal{A}_{\ell-2}$. These two matrices are block diagonal and have the structure prescribed by Theorem 11, so $A_1 A_2$ is block CMV.

On the other hand, every CMV matrix can be factored as $A = A_1 A_2$, and the diagonal blocks of these two matrices have the structure required to be the active part of a block unitary transformation. Therefore, both A_1 and A_2 can be written as the product of a sequence of odd and even-indexed block unitary transformations, respectively. This concludes the proof. \square

Remark 16. If we remove the assumption about the invertibility of the upper triangular blocks R, L^H of the unitary transformations \mathcal{A}_j in the previous lemma, then the factored representation still implies a block CMV shape of the cumulative matrix A . Indeed, these block CMV shaped matrices are the ones considered in the actual reduction process (compare also with Lemma 18).

LEMMA 17. Let $\mathcal{A}_1, \dots, \mathcal{A}_{\ell-1} \in \mathbb{C}^{n \times n}$ and $\mathcal{B}_{\ell-1}, \dots, \mathcal{B}_1 \in \mathbb{C}^{n \times n}$, $n = \ell k$, be two sequences of block unitary transformations of block size k . Then, there exist an $n \times n$ unitary matrix P and a sequence of block unitary transformations of block size k $\mathcal{C}_1, \dots, \mathcal{C}_{\ell-1} \in \mathbb{C}^{n \times n}$ such that

$$P \mathcal{A}_1 \dots \mathcal{A}_{\ell-1} \mathcal{B}_{\ell-1} \dots \mathcal{B}_1 P^H = \mathcal{C}_1 \dots \mathcal{C}_{\ell-1}.$$

Moreover, $P(e_1 \otimes I_k) = e_1 \otimes I_k$.

Proof. We prove the result by induction on ℓ . If $\ell = 1$, there is nothing to prove, and if $\ell = 2$, we can choose $P = I$ and we have

$$\mathcal{A}_1 \mathcal{B}_1 = \mathcal{C}_1,$$

which can be satisfied thanks to property (ii) of Lemma 14. Assume now that the result is valid for $\ell - 1$, and we want to prove it for ℓ . We can write

$$A := \mathcal{A}_1 \dots \mathcal{A}_{\ell-1} \mathcal{B}_{\ell-1} \dots \mathcal{B}_1 = \mathcal{A}_1 A_2 \mathcal{B}_1,$$

where we have set $A_2 := \mathcal{A}_2 \dots \mathcal{A}_{\ell-1} \mathcal{B}_{\ell-1} \dots \mathcal{B}_2$. Since the first block column and block row of A_2 is equal to the identity, we can rewrite it as $A_2 = I_k \oplus \tilde{A}$ with $\tilde{A} = \tilde{\mathcal{A}}_1 \dots \tilde{\mathcal{A}}_{\ell-2} \tilde{\mathcal{B}}_{\ell-2} \dots \tilde{\mathcal{B}}_1$. By the inductive hypothesis we have that there exists a unitary matrix \tilde{P} so that $\tilde{P} \tilde{A} \tilde{P}^H = \tilde{\mathcal{C}}_1 \dots \tilde{\mathcal{C}}_{\ell-2}$.

If we set $P_2 := I_k \oplus \tilde{P}$, then $P_2 \mathcal{A}_1 = \mathcal{A}_1$ and $P_2 \mathcal{B}_1^H = \mathcal{B}_1^H$, so we have

$$(2) \quad P_2 A P_2^H = \mathcal{A}_1 \mathcal{D}_2 \dots \mathcal{D}_{\ell-1} \mathcal{B}_1,$$

where we have set $\mathcal{D}_{j+1} := I_k \oplus \tilde{\mathcal{C}}_j$. We can move \mathcal{B}_1 to the left since it commutes with all the matrices except the first two, and by setting $\mathcal{A}_1 \mathcal{D}_2 \mathcal{B}_1 = \hat{\mathcal{B}}_2 \mathcal{C}_1 \hat{\mathcal{D}}_2$, thanks to property (iv) of Lemma 14, we get

$$P_2 A P_2^H = \hat{\mathcal{B}}_2 \mathcal{C}_1 \hat{\mathcal{D}}_2 \mathcal{D}_3 \dots \mathcal{D}_{\ell-1}.$$

Left-multiplying by $\hat{\mathcal{B}}_2^H$ and right-multiplying by $\hat{\mathcal{B}}_2$ yields

$$(\hat{\mathcal{B}}_2 P_2) A (\hat{\mathcal{B}}_2 P_2)^H = \mathcal{C}_1 \hat{\mathcal{D}}_2 \mathcal{D}_3 \dots \mathcal{D}_{\ell-1} \hat{\mathcal{B}}_2.$$

The part of the right-hand side that follows \mathcal{C}_1 , that is, the matrix $\hat{\mathcal{D}}_2 \hat{\mathcal{D}}_3 \dots \mathcal{D}_{\ell-1} \hat{\mathcal{B}}_2$, has the same form of the right-hand side in (2) but with one term less. We can reuse the same idea $\ell - 3$ times and obtain the desired decomposition. \square

We show here the special form required by the block unitary transformation is not very restrictive. In particular, every time that we have a CMV-like factorization for a matrix in terms of block unitary transformations, we can always perform a unitary transformation to obtain the required triangular structure inside the blocks.

LEMMA 18. *Let $A \in \mathbb{C}^{n \times n}$, $n = \ell k$, a unitary matrix that can be factored as*

$$A = \begin{cases} \mathcal{A}_1 \mathcal{A}_3 \dots \mathcal{A}_{\ell-1} \mathcal{A}_2 \mathcal{A}_4 \dots \mathcal{A}_{\ell-2} & \text{if } \ell \text{ is even,} \\ \mathcal{A}_1 \mathcal{A}_3 \dots \mathcal{A}_{\ell-2} \mathcal{A}_2 \mathcal{A}_4 \dots \mathcal{A}_{\ell-1} & \text{if } \ell \text{ is odd} \end{cases}$$

with $\mathcal{A}_j \in \mathbb{C}^{n \times n}$ block unitary transformation of block size k . Then, there exist a unitary transformation P , which is the direct sum of ℓ unitary blocks of size $k \times k$, such that $P A P^H = \tilde{A}$ satisfies

$$\tilde{A} = \begin{cases} \tilde{\mathcal{A}}_1 \tilde{\mathcal{A}}_3 \dots \tilde{\mathcal{A}}_{\ell-1} \tilde{\mathcal{A}}_2 \tilde{\mathcal{A}}_4 \dots \tilde{\mathcal{A}}_{\ell-2} & \text{if } \ell \text{ is even,} \\ \tilde{\mathcal{A}}_1 \tilde{\mathcal{A}}_3 \dots \tilde{\mathcal{A}}_{\ell-2} \tilde{\mathcal{A}}_2 \tilde{\mathcal{A}}_4 \dots \tilde{\mathcal{A}}_{\ell-1} & \text{if } \ell \text{ is odd,} \end{cases}$$

where

$$\tilde{\mathcal{A}}_j = \begin{bmatrix} \times & L \\ R & \times \end{bmatrix}, \quad R, L^H \text{ upper triangular.}$$

Proof. Assume that we are in the case ℓ even, and let us denote $A = A_l A_r = (\mathcal{A}_1 \mathcal{A}_3 \dots \mathcal{A}_{\ell-1}) \cdot (\mathcal{A}_2 \mathcal{A}_4 \dots \mathcal{A}_{\ell-2})$. It is not restrictive to assume that the blocks in A_r are already in the required form. In fact, if this is not true, we can compute a block diagonal unitary matrix Q (with $k \times k$ blocks) such that $Q A_r$ has the required shape. Then, by replacing A_l with $A_l Q^H$ we get another factorization of A where the right factor is already in the correct form.

We now show that we can take the left factor in the same shape without deteriorating the structure of the right one. Let Q_1 be a unitary transformation operating on the first block row such that $Q_1 A_l$ has the block in position (1, 2) in lower triangular form. Then we have

$$Q_1 A_l A_r Q_1^H = Q_1 A_l Q_1^H A_r = A_l^{(1)} A_r$$

since A_r and Q_1 commute. Moreover, $A_l^{(1)}$ has the first block row with the correct structure.

We now compute another unitary transformation Q_2 operating on the second block row such that $Q_2 A_l^{(1)}$ has the second row with the correct structure. Now the matrix $A_r Q_2^H$ loses the triangular structure in the block in position (3, 2). However,

we can compute another transformation P_3 operating on the third block row that restores the structure in $P_3 A_r Q_2^H$, and therefore we have

$$Q_2 A_l^{(1)} A_r Q_2^H = \left(Q_2 A_l^{(1)} P_3^H \right) \cdot (P_3 A_r Q_2^H) = A_l^{(2)} A_r^{(2)}$$

since the right multiplication by P_3^H does not degrade the partial structure that we have in $A_l^{(1)}$. We can iterate this process until all the blocks have the required structure, and this proves the lemma. \square

In what follows we refer to the matrix \tilde{A} defined in the previous lemma as a CMV structured unitary matrix with block size k even if there is no assumption about the invertibility of the triangular blocks. We now have all the tools required to perform the initial reduction of U to upper triangular form and of D to block CMV structure.

THEOREM 19. *Let $D \in \mathbb{C}^{n \times n}$ be a unitary diagonal matrix and $U \in \mathbb{C}^{n \times k}$ with $n = \ell k$ for some $\ell \in \mathbb{N}$. Then, there exists a unitary matrix P such that PDP^H is CMV structured with block size k and $PU(e_1 \otimes I_k) = (e_1 \otimes I_k)U_1$.*

Proof. The proof is divided in two stages. First, we show that we can build a unitary matrix P_U such that $P_U U = R$ is upper triangular. Then, we use Lemma 17 to construct another unitary matrix $P_C = I_k \oplus \tilde{P}_C$ so that $P_C \cdot (P_U DP_U^H) \cdot P_C^H$ is in block CMV shape. We then set $P := P_C P_U$ and conclude the proof.

In view of Lemma 13 we can set $P_U = \mathcal{A}_1 \dots \mathcal{A}_{\ell-1}$ so that $P_U U = R$. Applying the same transformation to D yields

$$P_U DP_U^H = \mathcal{A}_1 \dots \mathcal{A}_{\ell-1} D \mathcal{A}_{\ell-1}^H \dots \mathcal{A}_1^H = \mathcal{A}_1 \dots \mathcal{A}_{\ell-1} \mathcal{B}_{\ell-1} \dots \mathcal{B}_1,$$

where we have used property (ii) of Lemma 14 to merge D into the right block unitary transformations. We may now use Lemma 17 to obtain unitary transformations $\mathcal{C}_1, \dots, \mathcal{C}_{\ell-1}$ and a unitary matrix P_V so that

$$P_V \cdot (P_U DP_U^H) \cdot P_V^H = \mathcal{C}_1 \dots \mathcal{C}_{\ell-1}.$$

We now want to manipulate the above factorization to obtain one in CMV form, according to Lemma 15. Let us left-multiply the above by $\mathcal{C}_2 \dots \mathcal{C}_{\ell-1}$ and right-multiply it by its inverse. We get

$$(\mathcal{C}_2 \dots \mathcal{C}_{\ell-1}) \mathcal{C}_1 \dots \mathcal{C}_{\ell-1} (\mathcal{C}_2 \dots \mathcal{C}_{\ell-1})^H = \mathcal{C}_2 \dots \mathcal{C}_{\ell-1} \mathcal{C}_1.$$

Since \mathcal{C}_1 commutes with all the factors except \mathcal{C}_2 , the above expression is equal to

$$\mathcal{C}_2 \mathcal{C}_1 \mathcal{C}_3 \dots \mathcal{C}_{\ell-1}.$$

We can repeat the operation and move $\mathcal{C}_4 \dots \mathcal{C}_{\ell-1}$ to the left and, using the commutation properties, obtain

$$(\mathcal{C}_4 \dots \mathcal{C}_{\ell-1}) \mathcal{C}_2 \mathcal{C}_1 \mathcal{C}_3 \dots \mathcal{C}_{\ell-1} (\mathcal{C}_4 \dots \mathcal{C}_{\ell-1})^H = \mathcal{C}_4 \dots \mathcal{C}_{\ell-1} \mathcal{C}_2 \mathcal{C}_1 \mathcal{C}_3 = \mathcal{C}_2 \mathcal{C}_4 \mathcal{C}_1 \mathcal{C}_3 \mathcal{C}_5 \dots \mathcal{C}_{\ell-1}.$$

We repeat this process until we obtain all the \mathcal{C}_i with even indices on the left and the ones with odd indices on the right. Since this is the structure required by Lemma 15 and we have performed operations that do not involve \mathcal{C}_1 , we can write this final step as a matrix P_S with $P_S(e_1 \otimes I_k) = e_1 \otimes I_k$, so that

$$P_S P_V P_U A (P_S P_V P_U)^H = \mathcal{C}_2 \dots \mathcal{C}_{\ell-2} \mathcal{C}_1 \dots \mathcal{C}_{\ell-1},$$

assuming ℓ is even. Since $P_S P_V(e_1 \otimes I_k) = e_1 \otimes I_k$ by construction we have that $P_S P_V P_U U = R$. We can compute another unitary block diagonal transformation P_T obtained through Lemma 18 to ensure the correct triangular structure in the unitary blocks, and so the proof is complete by settings $P := P_T P_S P_V P_U$. \square

We now comment on the computational cost of performing the algorithm in the proof of Theorem 19. Assume that the matrices are of size n and the blocks are of size k . The following steps are required:

- (i) A block QR reduction of the matrix U . This requires the computation of $\mathcal{O}(\frac{n}{k})$ QR factorizations of $2k \times 2k$ matrices. The total cost is $\mathcal{O}(\frac{n}{k} \cdot k^3) \sim \mathcal{O}(nk^2)$.
- (ii) The application of Lemma 17, which requires about $\mathcal{O}((\frac{n}{k})^2)$ QR factorizations and matrix multiplications, all involving $2k \times 2k$ or $3k \times 3k$ matrices. Therefore, the total cost is of $\mathcal{O}(n^2k)$ flops.
- (iii) The manipulation of the order of the transformations requires again $\mathcal{O}((\frac{n}{k})^2)$ operations on $2k \times 2k$ matrices, thus adding another $\mathcal{O}(n^2k)$ flop count.
- (iv) The final application of Lemma 18 requires the computation of $\mathcal{O}(\frac{n}{k})$ QR factorizations, and thus another $\mathcal{O}(nk^2)$ contribution.


The total cost of the above steps is thus of $\mathcal{O}(n^2k)$ flops, since $k \leq n$, and the dominant part is given by steps (ii) and (iii).

3.6. Preserving the CMV structure. In this subsection we deal with efficiently implementing the final part of the algorithm: the reduction from block CMV form to upper Hessenberg form. In order to describe this part of the algorithm we consider the factored form that we have presented in the previous section, that is, the factorization (here reported for ℓ even)

$$A_{CMV} = \mathcal{A}_1 \dots \mathcal{A}_{\ell-1} \mathcal{B}_2 \dots \mathcal{B}_{\ell-2}.$$

In the following we are interested in showing how, using an appropriate bulge-chasing strategy, we can maintain the sparse structure of A_{CMV} during the Hessenberg reduction procedure of $A_{CMV} + UV^H$.

Given that the block structure and the indexing inside the blocks can be difficult to visualize, we complement the analysis with a pictorial description of the steps, in which the matrix has the following structure:




Let P_1 be a sequence of $2k - 2$ Givens rotations that operate on the rows with indices $2, \dots, 2k$, starting from the bottom to the top. These are the rotations needed to put the first column of $A_{CMV} + UV^H$ in upper Hessenberg form. Since we do not make any assumption on the structure of U and V beyond having U in upper triangular form (the block U_1 in Theorem 19 can be drained into the factor V), we do not restrict the analysis to the case where these rotations create some zeros in A_{CMV} .


Our task is to determine a set of unitary transformations P_2 , which only operate on rows and columns with indices larger than 3 (in order to not destroy the upper

Hessenberg structure created by P_1), such that $P_2 P_1 A_{CMV} P_1^H P_2^H$ has the trailing submatrix obtained removing the first row and column still in block CMV shape.


We start by applying the first $k - 1$ rotations in P_1 from both sides. Applying them from the left degrades the upper triangular structure of the block in position $(2, 1)$ and makes it upper Hessenberg. Pictorially, we obtain the following:



where $P_{1,1}$ has been used to denote the first $k - 1$ rotations in P_1 , and the gray area marks the fill-in introduced in the matrix. Applying them from the right, instead, creates some bulges in the upper triangular blocks in position $(4, 2)$. These bulges correspond with a rank one structure in the lower triangular part of this block and can be chased away by a sequence of $k - 1$ rotations. We can chase them by multiplying by these rotations from the left and from the right and then proceed in the same way until they reach the bottom of the matrix as follows:




We now continue to apply the rotations from the left, and this will create bulges in the lower triangular structure of the block $(1, 3)$. Similarly, these bulges have a rank one structure, and they can be chased until the bottom of the matrix, as shown in the next picture:



and we continue the procedure until we absorb the bulge at the bottom of the matrix. Handling the remaining rotations from the right is not so easy. Applying them at this stage would alter the first column of the block in position $(3, 2)$ and create a


rank 1 block in position (3, 1). Since we want to avoid this fill-in we roll-up the first column of the block (3, 2) by using some rotations on the left. As a consequence of the nullity theorem, this will automatically annihilate the diagonal entries of the lower triangular matrix in position (3, 5) and, at the last step, the first row of the block in position (3, 4). We perform a similar transformation (operating only from the left) on the other lower blocks. We obtain the following structure:



After this reduction we still have to perform some transformations from the right. These transformations only operate on odd-indexed blocks and thus commute with transformations operating on even-indexed blocks. This will be important in the following, since we keep these transformation “in standby,” and we will only apply them later.

We can now finally apply the transformations on the first block column. The first rotation will propagate the only nonzero element that we have left in the first column of the block (3, 2) and that will extend the upper Hessenberg structure of the block in position (2, 1). The following rotations will create bulges in the Hessenberg structure, which we will chase to the bottom of the matrix. Notice that we are able to do so since all the rotations operate on even-indexed blocks, thus commuting with the ones we have left in standby.

After the chasing procedure, we apply these last rotations. This creates some fill-in, as reported in the following picture.



We notice that the fill-in on the lower triangular matrices in the second block superdiagonal has reintroduced the elements that were annihilated due to the nullity theorem but shifted down and right by one. In fact, a careful look at the matrix reveals that the complete CMV structure has been shifted down and right by one entry, and thus removing the first row and column gives us a matrix with the same structure, as desired. We summarize the previous procedure in the following result.

THEOREM 20. *Let A_{CMV} a matrix in block CMV form, as defined previously. Let P a sequence of $2k - 2$ Givens rotations acting on the rows $2, \dots, 2k$, starting from the bottom to the top. Then there exist a unitary matrix Q such that*

$$\tilde{A} = QPA_{CMV}P^HQ^H$$

has the same block CMV structure of A_{CMV} in the trailing matrix obtained removing the first row and the first column. Moreover, QPU has only the first $k + 1$ rows different from zero.

Proof. Follow the steps described in the previous pages. The claim on the structure of U after the transformation follows by noticing the all the transformations except the ones contained in P operates on rows with indices larger than $k + 1$. \square

The same procedure can be applied in all the steps of Hessenberg reduction, and the CMV structure shifts down and right of one entry at each step. This provides a concrete procedure that allows one to perform the reduction to upper Hessenberg form in a cheap way.

We can perform a flops count for the above algorithm to ensure that the complexity is the one expected. First, notice that applying a Givens rotation to the matrix A has a cost of $\mathcal{O}(k)$ flops, given the banded structure that we are preserving. Following the algorithm in Theorem 20 shows that we only need $\mathcal{O}(n)$ rotations to perform a single step, and therefore we have $\mathcal{O}(nk)$ flops per step, which yields a total cost of $\mathcal{O}(n^2k)$ operations.

However, while the algorithm allows one to compute the Givens rotations needed at each step in order to continue the process, it is not clear, at this stage, how to recover the reduced matrix at the end of the procedure. This is the topic of the next section.

Remark 21. Notice that, after the first step, the transformed matrix \tilde{A} has the trailing part in block CMV form, but the trailing matrix is not unitary. Therefore, in order to apply the algorithm in the following steps (and still make use of the nullity theorem), it is necessary to formally consider the entire matrix and just limit the chasing on the trailing part. Also notice that in the proof of Theorem 20 we rely on the nullity theorem to guarantee that some elements have to be zero. This does not happen in practice due to floating-point arithmetic, but since everything here is operating on a unitary matrix after some steps we have $\text{fl}(QA_{CMV}Q^H) = QA_{CMV}Q^H + \Delta$ and the entries in the perturbation have to be $\approx u$ normwise without error amplification.

3.7. Representing the final matrix. The previous section provides a concrete way to compute the transformations required to take the diagonal matrix plus low-rank correction into upper Hessenberg form. The computation of the transformations is possible since, at each step, we have a structured representations of the part of the matrix that needs to be reduced. However, with the previous approach, we do not have an explicit expression of the part of the matrix that has already been reduced to upper Hessenberg form. The aim of this section is to show how to represent that part.

To perform this analysis, we use the following notation: We denote by P_j and by Q_j , respectively, the product of $2k - 2$ rotations that cleans the elements in the j th column in order to make the matrix upper Hessenberg in that column and the product of rotations applied to perform the subsequent bulge chasing. Thus the matrix after j steps is $Q_j P_j \dots Q_1 P_1 (A_{CMV} + UV^H) P_1^H Q_1^H \dots P_j^H Q_j^H$, and it will continue to have a $2k$ -banded structure in the bottom-right part that still have to be reduced.

Let then $PA_{CMV}P^H$ be the the unitary matrix obtained at the end of the process, that is, $P = Q_{n-2}P_{n-2} \dots Q_1P_1$.

LEMMA 22. *Each row of $PA_{CMV}P^H$ admits a representation of the form*

$$e_j^H PA_{CMV}P^H = [\mathbf{w}^H \ 0 \ \dots \ 0] P_{j+1}^H \dots P_{n-2}^H,$$

where $\mathbf{w} \in \mathbb{C}^{j+2k}$ for any $j \leq n - 2k$.

Proof. It follows directly by the preservation of the CMV structure (and so in particular of the $2k$ -banded structure) and by noting that after reducing the j th column the rotations on the left will not alter the j th row anymore, and the ones on the right will update it only when performing the first cleaning step, but not the bulge chasing (since these will act on zero components). \square

Notice that a similar result also holds for the columns, since the situation is completely symmetric. The above representation is usually known as Givens–Vector representation [24, 25] and can also be seen as a particular case of generators-based representation [15, 16]. Being the matrix upper Hessenberg we only need to store the upper triangular part plus the subdiagonal, so the above representation needs only $O(nk)$ storage, since we need to represent $O(n)$ vectors of length k and $O(nk)$ rotations.

4. Numerical results. In this section we report numerical experiments that validate the proposed algorithm. We tested the accuracy and the runtime of our implementations.

The software can be downloaded for further testing at <http://numpi.dm.unipi.it/software/rank-structured-hessenberg-reduction>. The package contains two functions:

- `rshr_d1r` implements the Hessenberg reduction for a real diagonal plus low rank matrix. This function is implemented in MATLAB and FORTRAN, and a MATLAB wrapper to the FORTRAN code is provided and used automatically if possible.
- `rshs_u1r` implements the analogous reduction for the unitary plus low-rank case. It is only implemented as a MATLAB function.

4.1. The real diagonal plus low rank case. We have checked the complexity of our reduction strategy by fixing n and k , respectively, and comparing the runtime while varying the other parameter. We verified that the complexity is quadratic in n and linear in k .

When measuring the complexity in n we have also compared our timings with the DGEHRD function included in LAPACK 3.6.1.

The tests have been run on a server with an Intel Xeon CPU E5-2697 running at 2.60GHz and with 128 GB of RAM. The amount of RAM allowed us to test DGEHRD for large dimensions, which would have not been feasible on a desktop computer.

The results that we report are for compiler optimizations turned off. This makes the dependency of the timings on the size and rank much clearer in the plots and allows us to clearly connect the timings with the flop count. However, each experiment has been repeated with optimization turned on that yields a gain of a factor between 2 and 4 in the timings, both for DGEHRD and for our method, so there is no significant difference when comparing the approaches.

Figure 1 shows the complexity of our method as a function of the size n . The behavior is clearly quadratic, as shown by the dashed line in the plot. It is also clear that the method is faster than LAPACK for relatively small dimensions (in this case, $n \approx 16$) if the rank is small. Given that the ratio between the two complexities is $n^3/n^2k = \frac{n}{k}$ this suggests that our approach is faster when $k \lesssim \frac{n}{8}$. As an example, consider the timings for $k = 32$ that are also reported in Figure 1; we expect our approach to be faster when $n > 8k = 256$, and this is confirmed by the plot.

In particular, our constant in front of the complexity is just 8 times larger than a dense Hessenberg reduction. We believe that a clever use of blocking might even improve this result and could be the subject of future investigation.




FIG. 1. Complexity of the Hessenberg reduction method in dependency of the size n . The quasi-separable rank k has been chosen equal to 4 and 32 in this example. In each example the tests have been run only when $k < n$.




FIG. 2. Complexity of the Hessenberg reduction method in dependency of the quasi-separable rank k . The dimension n has been chosen equal to 2048 in this example, and rank varies from 1 to 256.

We performed a similar analysis for the cost as a function of the rank. Figure 2 shows that the growth of the complexity with respect to k is essentially linear. We have taken a matrix of size 2048 and have compared the performance of our own implementation with the one in LAPACK. It is worth noting that since our implementation only requires $\mathcal{O}(nk)$ storage if n grows large enough, then it also provides an alternative when a dense reduction cannot be applied due to memory constraints. Moreover, the linear behavior is very clear from the plot.

Our implementation relies on a compressed diagonal storage scheme in order to efficiently operate on the small bandwidth, and it never needs to store the large $n \times n$ matrix.




FIG. 3. Backward error on the original matrix of the computed Hessenberg form for various values of n and k in the real diagonal plus low rank case. The dashed line represents the machine precision multiplied by the square root of the size of the matrix.

Finally, we have tested the backward stability of our approach. In order to do so we have computed the Hessenberg form with our fast method, along with the change of basis Q by accumulating all the plane rotations. Then, we have measured the norm of $A - Q^H H Q$. The results are reported in Figure 3.

The experiments show that our method is backward stable with a backward error smaller than $\mathcal{O}(\sqrt{nu})$, where n is the size of the matrix and u the unit roundoff.

One can prove from Algorithm in section 2 that the backward stability is achieved with a low-degree polynomial in n and k times the unit roundoff, obtaining an upper bound of the form $\mathcal{O}(n^2ku)$. However, this bound is never reached in practice since the errors do not align all in the same direction. In fact, Figure 3 shows that the error grows as $\mathcal{O}(\sqrt{n} \cdot u)$. The LAPACK implementation DGEHRD manages to be even more accurate, and the growth seems to be very slow.

The numerical results have been obtained by averaging different runs of the algorithm on random matrices.

4.2. The unitary case. In the unitary diagonal plus low rank case we have used our MATLAB implementation to test the accuracy of the algorithm. We also report some results on the timings that show that the asymptotic complexity is the one we expect.

The implementation of the unitary plus low-rank case is slightly more involved than the Hermitian one, mainly due to the bookkeeping required to handle the application of plane rotations in the correct order. Nevertheless, the software developed for the real diagonal plus low rank case could be adapted to efficiently handle also this case.

We have reported the results for the timings (taken in MATLAB) for our implementation of the Hessenberg reduction of unitary plus low rank matrices in Figures 4 and 5. These two figures report the dependency on the size (which is quadratic) and on the quasi-separable rank k (which is linear).

We believe that an efficient implementation of this reduction could be faster than LAPACK for small dimension, as happens for the real diagonal plus low rank case.




FIG. 4. Complexity of the Hessenberg reduction method in dependency of the dimension n of the unitary plus low rank case. The dimension n has been chosen between 16 and 1024 with a fixed rank $k = 4$.




FIG. 5. Complexity of the Hessenberg reduction method in dependency of the quasi-separable rank k for the unitary plus low rank case. The rank k has been chosen between 16 and 512 with a fixed size $n = 8192$.

However, due to the increased complexity, it is likely that this method will be faster for slightly larger ratios of $\frac{n}{k}$, compared to what we have for the real diagonal plus low rank case.

Similarly to the Hermitian case, we have also performed numerical experiments with the purpose of estimating the backward error. The results are reported in Figure 6 and show a moderate growth of the backward error when the size increases.

5. Conclusions and future work. In this paper we have presented a fast algorithm for reducing a $n \times n$ real/unitary diagonal D plus a low rank correction in upper Hessenberg form. To our knowledge this is the first algorithm which is




FIG. 6. Backward error on the original matrix of the computed Hessenberg form for various values of n and k in the unitary plus low rank case. The dashed line represents the machine precision multiplied by the square root of the size of the matrix.

efficient w.r.t. the size k of the correction by requiring $O(n^2k)$ flops and provides a viable alternative to the dense LAPACK routine for small sizes. The approach for the unitary case relies upon some theoretical and computational properties of block unitary CMV matrices of independent interest for computations with orthogonal matrix polynomials. The application of our algorithm for solving interpolation-based linearizations of generalized eigenvalue problems is still ongoing research.

In particular, the output of both reduction algorithms is already in a rank structured form which could be exploited in an iteration (such as the QR method) to compute the eigenvalues of the matrix A . This could provide a method with optimal complexity for the computation of eigenvalues of matrix polynomials, expressed in monomial and interpolation bases.

REFERENCES

- [1] A. AMIRASLANI, R. M. CORLESS, AND P. LANCASTER, *Linearization of matrix polynomials expressed in polynomial bases*, IMA J. Numer. Anal., 29 (2009), pp. 141–157, <https://doi.org/10.1093/imanum/drm051>.
- [2] G. S. AMMAR AND W. B. GRAGG, *$O(n^2)$ reduction algorithms for the construction of a band matrix from spectral data*, SIAM J. Matrix Anal. Appl., 12 (1991), pp. 426–431, <https://doi.org/10.1137/0612030>.
- [3] P. ARBENZ AND G. H. GOLUB, *On the spectral decomposition of Hermitian matrices modified by low rank perturbations with applications*, SIAM J. Matrix Anal. Appl., 9 (1988), pp. 40–58, <https://doi.org/10.1137/0609004>.
- [4] Y. ARLINSKIĬ, *Conservative discrete time-invariant systems and block operator CMV matrices*, Methods Funct. Anal. Topology, 15 (2009), pp. 201–236.
- [5] R. BEVILACQUA, G. M. DEL CORSO, AND L. GEMIGNANI, *A CMV-based eigensolver for companion matrices*, SIAM J. Matrix Anal. Appl., 36 (2015), pp. 1046–1068, <https://doi.org/10.1137/140978065>.
- [6] R. BEVILACQUA, G. M. DEL CORSO, AND L. GEMIGNANI, *Compression of unitary rank-structured matrices to CMV-like shape with an application to polynomial rootfinding*, J. Comput. Appl. Math., 278 (2015), pp. 326–335, <https://doi.org/10.1016/j.cam.2014.09.023>.
- [7] D. A. BINI AND L. ROBOL, *On a class of matrix pencils and ℓ -ifications equivalent to a given matrix polynomial*, Linear Algebra Appl., 502 (2016), pp. 275–298, <https://doi.org/10.1016/j.laa.2015.07.017>.

- [8] D. A. BINI AND L. ROBOL, *Quasiseparable Hessenberg reduction of real diagonal plus low rank matrices and applications*, Linear Algebra Appl., 502 (2016), pp. 186–213, <https://doi.org/10.1016/j.laa.2015.08.026>.
- [9] C. H. BISCHOF, B. LANG, AND X. SUN, *A framework for symmetric band reduction*, ACM Trans. Math. Software, 26 (2000), pp. 581–601, <https://doi.org/10.1145/365723.365735>.
- [10] A. BUNSE-GERSTNER AND L. ELSNER, *Schur parameter pencils for the solution of the unitary eigenproblem*, Linear Algebra Appl., 154/156 (1991), pp. 741–778, [https://doi.org/10.1016/0024-3795\(91\)90402-1](https://doi.org/10.1016/0024-3795(91)90402-1).
- [11] M. J. CANTERO, L. MORAL, AND L. VELÁZQUEZ, *Five-diagonal matrices and zeros of orthogonal polynomials on the unit circle*, Linear Algebra Appl., 362 (2003), pp. 29–56, [https://doi.org/10.1016/S0024-3795\(02\)00457-3](https://doi.org/10.1016/S0024-3795(02)00457-3).
- [12] S. DELVAUX AND M. VAN BAREL, *A Hessenberg reduction algorithm for rank structured matrices*, SIAM J. Matrix Anal. Appl., 29 (2007), pp. 895–926, <https://doi.org/10.1137/060658953>.
- [13] Y. EIDELMAN, L. GEMIGNANI, AND I. GOHBERG, *Efficient eigenvalue computation for quasiseparable Hermitian matrices under low rank perturbations*, Numer. Algorithms, 47 (2008), pp. 253–273, <https://doi.org/10.1007/s11075-008-9172-0>.
- [14] Y. EIDELMAN, I. GOHBERG, AND L. GEMIGNANI, *On the fast reduction of a quasiseparable matrix to Hessenberg and tridiagonal forms*, Linear Algebra Appl., 420 (2007), pp. 86–101, <https://doi.org/10.1016/j.laa.2006.06.028>.
- [15] Y. EIDELMAN, I. GOHBERG, AND I. HAIMOVICI, *Separable Type Representations of Matrices and Fast Algorithms. Vol. 1: Basics. Completion Problems. Multiplication and Inversion Algorithms*, Oper. Theory Adv. Appl. 234, Birkhäuser/Springer, Basel, 2014.
- [16] Y. EIDELMAN, I. GOHBERG, AND I. HAIMOVICI, *Separable Type Representations of Matrices and Fast Algorithms. Vol. 2: Eigenvalue Method*, Oper. Theory Adv. Appl. 235, Birkhäuser/Springer, Basel, 2014.
- [17] M. FIEDLER AND T. L. MARKHAM, *Completing a matrix when certain entries of its inverse are specified*, Linear Algebra Appl., 74 (1986), pp. 225–237, [https://doi.org/10.1016/0024-3795\(86\)90125-4](https://doi.org/10.1016/0024-3795(86)90125-4).
- [18] I. GOHBERG, P. LANCASTER, AND L. RODMAN, *Matrix Polynomials*, Classics Appl. Math., Academic Press, New York, 1982.
- [19] W. H. GUSTAFSON, *A note on matrix inversion*, Linear Algebra Appl., 57 (1984), pp. 71–73, [https://doi.org/10.1016/0024-3795\(84\)90177-0](https://doi.org/10.1016/0024-3795(84)90177-0).
- [20] B. KÄGSTRÖM, D. KRESSNER, E. S. QUINTANA-ORTÍ, AND G. QUINTANA-ORTÍ, *Blocked algorithms for the reduction to Hessenberg-triangular form revisited*, BIT, 48 (2008), pp. 563–584, <https://doi.org/10.1007/s10543-008-0180-1>.
- [21] R. KILLIP AND I. NENCIU, *CMV: The unitary analogue of Jacobi matrices*, Comm. Pure Appl. Math., 60 (2007), pp. 1148–1188, <https://doi.org/10.1002/cpa.20160>.
- [22] B. SIMON, *CMV matrices: five years after*, J. Comput. Appl. Math., 208 (2007), pp. 120–154, <https://doi.org/10.1016/j.cam.2006.10.033>.
- [23] M. VAN BAREL AND A. BULTHEEL, *Orthonormal polynomial vectors and least squares approximation for a discrete inner product*, Electron. Trans. Numer. Anal., 3 (1995), pp. 1–23.
- [24] R. VANDEBRIL, M. VAN BAREL, AND N. MASTRONARDI, *Matrix Computations and Semiseparable Matrices. Vol. I: Linear Systems*, Johns Hopkins University Press, Baltimore, MD, 2008.
- [25] R. VANDEBRIL, M. VAN BAREL, AND N. MASTRONARDI, *Matrix Computations and Semiseparable Matrices. Vol. II: Eigenvalue and Singular Value Methods*, Johns Hopkins University Press, Baltimore, MD, 2008.