

COMPUTING FUNCTIONS OF SYMMETRIC HIERARCHICALLY SEMISEPARABLE MATRICES*

ANGELO A. CASULLI[†], DANIEL KRESSNER[‡], AND LEONARDO ROBOL[§]

Abstract. The aim of this work is to develop a fast algorithm for approximating the matrix function $f(A)$ of a square matrix A that is symmetric and has hierarchically semiseparable (HSS) structure. Appearing in a wide variety of applications, often in the context of discretized (fractional) differential and integral operators, HSS matrices have a number of attractive properties facilitating the development of fast algorithms. In this work, we use an unconventional telescopic decomposition of A , inspired by recent work of Levitt and Martinsson on approximating an HSS matrix from matrix-vector products with a few random vectors. This telescopic decomposition allows us to approximate $f(A)$ by recursively performing low-rank updates with rational Krylov subspaces while keeping the size of the matrices involved in the rational Krylov subspaces small. In particular, no large-scale linear system needs to be solved, which yields favorable complexity estimates and reduced execution times compared to existing methods, including an existing divide-and-conquer strategy. The advantages of our newly proposed algorithms are demonstrated for a number of examples from the literature, featuring the exponential, the inverse square root, and the sign function of a matrix. For the special case of matrix inversion, our algorithm reduces to a procedure previously proposed by Gillman, Young, and Martinsson [*Front. Math. China*, 7 (2012), pp. 217–247].

Key words. functions of matrices, hierarchically semiseparable, rational Krylov

MSC codes. 65F60, 15B99

DOI. 10.1137/24M1642354

1. Introduction. Consider a symmetric matrix $A \in \mathbb{R}^{n \times n}$ with spectral decomposition $A = V\Lambda V^T$, with the orthogonal matrix V and the diagonal matrix $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$ containing the eigenvalues of A . Given a scalar function f well defined on the eigenvalues of A , the matrix function $f(A) \in \mathbb{R}^{n \times n}$ is defined as $Vf(\Lambda)V^T$, where $f(\Lambda) := \text{diag}(f(\lambda_1), \dots, f(\lambda_n))$. Popular examples include the matrix inverse, the matrix exponential, the sign function, and the (inverse) matrix square root; see the monograph [16] for an overview. When A is of moderate size, $f(A)$ can simply be computed according to this definition, via computing the spectral decomposition of A , or using a more specialized algorithm such as the scaling-and-squaring method for the matrix exponential [17]. These methods typically require $\mathcal{O}(n^2)$ memory and $\mathcal{O}(n^3)$ operations, and thus become infeasible for larger n . If only the computation of $f(A)B$ for a (block) vector B is needed, (rational) Krylov subspace methods are well suited when A is large and (data) sparse; see [12] and the references therein.

The task of approximating the whole matrix function $f(A)$ for a large-scale matrix A is rather challenging and certainly requires additional assumptions on the data

*Received by the editors February 27, 2024; accepted for publication (in revised form) August 8, 2024; published electronically December 2, 2024.

<https://doi.org/10.1137/24M1642354>

Funding: The third author has been supported by the National Research Center in High Performance Computing, Big Data, and Quantum Computing (CN1 – Spoke 6) and by the PRIN 2022 Project “Low-Rank Structures and Numerical Methods in Matrix and Tensor Computations and Their Application.” The third author also acknowledges the MIUR Excellence Department Project awarded to the Department of Mathematics, University of Pisa, CUP I57G22000700001.

[†]Gran Sasso Science Institute, L’Aquila, Italy (angelo.casulli@gssi.it).

[‡]Institute of Mathematics, EPFL Lausanne, Lausanne, Switzerland (daniel.kressner@epfl.ch).

[§]Department of Mathematics, University of Pisa, Pisa, Italy (leonardo.robol@unipi.it).

sparsity structure of A . For example, if A is a banded matrix and f can be well approximated by a low-degree polynomial on the spectrum of A , then $f(A)$ can also be well approximated by a banded matrix [3], leading to fast algorithms, such as the ones described in [6, 8, 26]. If, on the other hand, f does not admit good polynomial approximations, then $f(A)$ usually does not admit a good approximation by a banded or, more generally, by a sparse matrix even when A is banded. Examples include the (inverse) square root or the sign function when A has eigenvalues that are close to zero relative to the width of the spectrum. For these examples, f still admits good *rational* approximations and the approximation of $f(A)$ can potentially be addressed using hierarchical low-rank techniques [14].

A matrix A is said to be *hierarchically off-diagonal low-rank* (HODLR) if it can be recursively block partitioned in a matrix with low-rank off-diagonal blocks, more specifically, there exists a block partitioning

$$A = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix}$$

such that A_{12} and A_{21} are of low rank and A_{11} and A_{22} are square matrices that can be recursively partitioned in the same way (until a minimal size of the diagonal blocks is reached). *Hierarchically semiseparable* (HSS) matrices additionally impose that the low-rank factors representing the off-diagonal blocks on the different levels of the recursion are nested; see section 2 for the precise definition. Hierarchical matrices, such as HODLR and HSS, admit data-sparse representations and cover a wide variety of matrix structures, including banded matrices and rational functions thereof. In particular, they can be used to approximate $f(A)$ when A is banded or, more generally, HSS whenever f admits a good rational approximation [6]. This property has been exploited to develop iterate-and-truncate methods using hierarchical matrices [11, 19] as well as a divide-and-conquer procedure based on low-rank updates and rational Krylov subspaces [6]. While numerical experiments in [6] show that the latter method is often preferable in terms of efficiency, it only exploits HODLR structure even when the matrix is HSS, and therefore does not fully benefit from the nestedness of low-rank factors in the HSS format.

To fully benefit from HSS structure, we will represent an HSS matrix in an unconventional way, via a telescopic decomposition. Such a decomposition was recently used by Levitt and Martinsson [20] to compute an HSS approximation of a matrix A from a few random matrix-vector products.¹ A symmetric matrix A in (standard) telescopic decomposition of depth L takes the form


$$A = D + U A^{(L-1)} U^T$$

with the block diagonal matrices

$$D = \begin{bmatrix} A_{11} & & & \\ & A_{22} & & \\ & & \ddots & \\ & & & A_{2^L, 2^L} \end{bmatrix}, \quad U = \begin{bmatrix} U_1 & & & \\ & U_2 & & \\ & & \ddots & \\ & & & U_{2^L} \end{bmatrix},$$

and $A^{(L-1)}$ in telescopic decomposition of depth $L-1$; see Figure 1 for an illustration of $L=2$. The number of columns of each U_i is determined by the HSS rank of A .

¹It is worth mentioning that another method for approximating HSS matrices from random matrix-vector products was recently presented by Halikias and Townsend [15], but this algorithm is not based on the telescopic decompositions considered in this work.

FIG. 1. *Telescopic decomposition of depth 2.*

Our approach is based on two observations: (1) The evaluation of $f(\mathbf{D})$ reduces to computing the small matrix functions $f(A_{ii})$ for $i = 1, \dots, 2^L$. (2) For small HSS rank, the difference between $A - \mathbf{D}$ has low rank and thus the low-rank update techniques from [1] apply. This allows us to (very accurately) approximate $f(A)$ as follows:

$$(1.1) \quad f(A) \approx f(\mathbf{D}) + \mathbf{W} \left(f(\mathbf{W}^T (\mathbf{D} + \mathbf{U} A^{(L-1)} \mathbf{U}^T) \mathbf{W}) - f(\mathbf{W}^T \mathbf{D} \mathbf{W}) \right) \mathbf{W}^T,$$

where $\mathbf{W}$ is the orthonormal basis of a rational Krylov subspace for the matrix $\mathbf{D}$ and starting block vector $\mathbf{U}$. As we will see in section 4, $\mathbf{W}$ inherits the block diagonal structure from $\mathbf{D}$, $\mathbf{U}$. Thus, $f(\mathbf{W}^T \mathbf{D} \mathbf{W})$ again reduces to applying f to the diagonal blocks of $\mathbf{W}^T \mathbf{D} \mathbf{W}$. More importantly, $\mathbf{W}^T (\mathbf{D} + \mathbf{U} A^{(L-1)} \mathbf{U}^T) \mathbf{W}$ is again in telescopic decomposition, which allows us to approximate the first term of (1.1) by recursively repeating the described procedure. The approximation to $f(A)$ eventually returned by the whole procedure is in a telescopic decomposition that differs from the ones in [20] and [21]. This will require us to introduce a more general notion of telescopic decompositions in this work.

Compared to the divide-and-conquer method described in [6], our algorithm fully exploits the nestedness relations of the HSS format. In particular, our construction only involves rational Krylov subspaces for block diagonal matrices instead of HSS matrices. This translates into reduced complexity and, in many cases, significantly lower execution times. Based on telescopic decompositions, our new algorithm combines well with the method by Levitt and Martinsson. This combination allows one to extract an approximation to $f(A)$ from the product of A with a few random vectors.

The remainder of the paper is structured as follows: The usual data-sparse representation of HSS matrices is described in section 2, while section 3 introduces a general concept of telescopic decompositions and shows how to convert between different forms of telescopic decompositions and the usual data-sparse representation of HSS matrices. In section 4, our newly proposed algorithm for approximating matrix functions is described and its convergence is analyzed. Section 5 recalls existing rational approximation results, which yield favorable complexity estimates for our algorithm in important special cases. This translates into superior efficiency, as demonstrated with several numerical examples in section 6.

Note that for most of sections 2 and 3, we will not assume that A is symmetric. Imposing symmetry would not simplify the exposition and our considerations on telescopic decompositions in section 3 might be of independent interest. In section 3.4, we will discuss the consequences of A being symmetric.

2. Hierarchically semiseparable matrices. To define hierarchically semiseparable (HSS) matrices we need to introduce a way to recursively split row and column indices of a matrix. Given a vector of indices $I = [1, 2, \dots, n]$, we use a perfect binary

tree $\mathcal{T}$, called *cluster tree*, to define subsets of indices obtained by subdividing I . The root γ of the tree is associated with the full vector I ; the rest of the tree is recursively defined in the following way: given a nonleaf node τ associated with an index vector I_τ , its children α, β , are associated with two vectors of consecutive indices I_α, I_β , such that I_τ is the concatenation of I_α and I_β . The depth of a node is defined as the distance from the root of the tree. The depth of the tree is denoted by L . We observe that $\mathcal{T}$ is uniquely defined by the index vectors associated with the nodes, and, hence, cluster trees can be defined by simply specifying the indices associated with the leaf nodes. For each leaf node α , we use $|\alpha|$ to denote the length of the vector of indices associated with α . We assume that the leaf nodes are “small,” that is, for some prescribed *threshold size* t it holds that $|\alpha| \leq t$ for every leaf α .

Given a matrix $A \in \mathbb{R}^{n \times n}$, we let $A_{\tau, \tau'}$ denote the submatrix of A obtained by selecting the row and column indices associated with nodes τ and τ' , respectively. In particular, we will consider diagonal blocks ($\tau = \tau'$) and sub/supdiagonal blocks (τ and τ' are siblings, i.e., children of the same node). We are now ready to state the definition of an HSS matrix following [21, section 3.3].

DEFINITION 2.1. *Given a cluster tree $\mathcal{T}$ for the indices $[1, \dots, n]$, a matrix $A \in \mathbb{R}^{n \times n}$ is called an HSS matrix of HSS rank r if the following hold:*

1. *For each pair of sibling nodes $\tau, \tau' \in \mathcal{T}$, there exist matrices $U_\tau^{(\text{big})} \in \mathbb{R}^{|\tau| \times r}$, $V_{\tau'}^{(\text{big})} \in \mathbb{R}^{|\tau'| \times r}$ with orthonormal columns and $\tilde{A}_{\tau, \tau'} \in \mathbb{R}^{r \times r}$, such that*

$$A_{\tau, \tau'} = U_\tau^{(\text{big})} \tilde{A}_{\tau, \tau'} (V_{\tau'}^{(\text{big})})^T.$$

2. *For each nonleaf node $\tau \in \mathcal{T}$ with children α and β , there exist $U_\tau, V_\tau \in \mathbb{R}^{2r \times r}$ with orthonormal columns, such that*

$$(2.1) \quad U_\tau^{(\text{big})} = \begin{bmatrix} U_\alpha^{(\text{big})} & \\ & U_\beta^{(\text{big})} \end{bmatrix} U_\tau \quad \text{and} \quad V_\tau^{(\text{big})} = \begin{bmatrix} V_\alpha^{(\text{big})} & \\ & V_\beta^{(\text{big})} \end{bmatrix} V_\tau.$$

Remark 2.2. For the purpose of simplifying the description, Point 1 of Definition 2.1 imposes the same rank r on every sub/supdiagonal block $A_{\tau, \tau'}$. All definitions and algorithms discussed in this work can be adapted to accommodate different ranks for different τ, τ' . In particular, our software implementation allows for nonconstant ranks in order to optimize storage and computational cost.

Definition 2.1 corresponds to the usual way of storing an HSS matrix A [28, 23]. For each leaf node τ , the matrices $U_\tau := U_\tau^{(\text{big})}$, $V_\tau := V_\tau^{(\text{big})}$ and for each nonleaf node τ the matrices U_τ, V_τ from (2.1) are stored. The latter pair of matrices are usually called *translation operators*. The nestedness relation (2.1) allows us to recursively recover $U_\alpha^{(\text{big})}$ and $V_\alpha^{(\text{big})}$ for any $\alpha \in \mathcal{T}$ from the (small) matrices U_τ, V_τ . We only need to additionally store the $r \times r$ matrices $\tilde{A}_{\tau, \tau'}$ for every pair of sibling nodes τ, τ' and the small diagonal blocks $A_{\tau, \tau}$ (of size at most t) for every leaf τ in order to recover the whole matrix A . To see this, let α and β denote the children of the root γ . Then

$$(2.2) \quad A = A_{\gamma, \gamma} = \begin{bmatrix} A_{\alpha, \alpha} & U_\alpha^{(\text{big})} \tilde{A}_{\alpha, \beta} (V_\beta^{(\text{big})})^T \\ U_\beta^{(\text{big})} \tilde{A}_{\beta, \alpha} (V_\alpha^{(\text{big})})^T & A_{\beta, \beta} \end{bmatrix}.$$

If α and β are nonleaf nodes, the matrices $A_{\alpha, \alpha}$ and $A_{\beta, \beta}$ can be recursively recovered in the same fashion. Otherwise, if α and β are leaves, these matrices are stored explicitly. In summary, the matrices

$$(2.3) \quad \{U_\tau, V_\tau : \tau \in \mathcal{T}\}, \quad \{\tilde{A}_{\tau, \tau'} : \tau, \tau' \text{ sibling nodes}\}, \quad \{A_{\alpha, \alpha} : \alpha \text{ leaf node}\}$$

define a data-sparse representation of an HSS matrix A .

Remark 2.3. There is no need to impose orthonormality on U_τ, V_τ in the representation (2.3). The orthogonality properties required by Definition 2.1 can always be ensured by the orthogonalization procedure described in [21, section 4.2], without changing the matrix A represented by (2.3) through the recursion (2.2).

2.1. Block diagonal matrices and depth reduction. To simplify notation, we make use of the following form of block diagonal matrices. Let $\{M_\tau\}_{\tau \in \mathcal{T}}$ be a set of matrices for a cluster tree $\mathcal{T}$. Then

$$(2.4) \quad \mathbf{M}^{(\ell)} := \text{blkdiag}(M_\tau : \tau \in \mathcal{T}, \text{depth}(\tau) = \ell)$$

denotes the block diagonal matrix with the diagonal blocks consisting of all matrices M_τ for which τ has depth equal to ℓ .

For example, given the data-sparse representation (2.3) of an HSS matrix A , the matrices $\mathbf{U}^{(L)}$ and $\mathbf{V}^{(L)}$ are block diagonal matrices containing the orthonormal matrices U_α and V_α , respectively, for every leaf $\alpha \in \mathcal{T}$ as diagonal blocks. When considering the product

$$\hat{A} := (\mathbf{U}^{(L)})^T A \mathbf{V}^{(L)},$$

the representation (2.3) is affected as follows:

$$(2.5) \quad \begin{aligned} &\{U_\tau, V_\tau : \text{depth}(\tau) \leq L-1\}, \quad \{U_\alpha^T U_\alpha, V_\alpha^T V_\alpha : \alpha \text{ leaf node}\}, \\ &\{\tilde{A}_{\tau, \tau'} : \tau, \tau' \text{ sibling nodes}\}, \quad \{U_\alpha^T A_{\alpha, \alpha} V_\alpha : \alpha \text{ leaf node}\}. \end{aligned}$$

This representation allows us to reconstruct the matrix $\hat{A}$ by a recursion analogous to (2.2). Since $U_\alpha^T U_\alpha = I$ and $V_\alpha^T V_\alpha = I$, the orthogonality properties of Definition 2.1 are satisfied as well. However, the size of $\hat{A}$ is $2^L m$, which is generally different from n and requires the cluster tree $\mathcal{T}$ to be adjusted accordingly.

DEFINITION 2.4. Given integers m and L , consider the index vector $[1, 2, \dots, 2^L m]$. Then $\mathcal{T}_m^{(L)}$ denotes the corresponding balanced cluster tree of depth L , having leaves associated with $[(i-1)m+1, \dots, im]$ for $i = 1, \dots, 2^L$.

In particular, the matrix $\hat{A}$ defined above is an HSS matrix associated with the cluster tree $\mathcal{T}_m^{(L)}$. When dropping the leaves, one obtains the balanced cluster tree $\mathcal{T}_{2m}^{(L-1)}$ of depth $L-1$. The matrix $\hat{A}$ is still an HSS matrix associated with $\mathcal{T}_{2m}^{(L-1)}$ but its parametrization (2.5) reduces to

$$\{U_\tau, V_\tau : \tau \in \mathcal{T}_{2m}^{(L-1)}\}, \quad \{\tilde{A}_{\tau, \tau'} : \tau, \tau' \text{ siblings in } \mathcal{T}_{2m}^{(L-1)}\}, \quad \{\hat{A}_{\alpha, \alpha} : \alpha \text{ leaf in } \mathcal{T}_{2m}^{(L-1)}\},$$

which matches the form of (2.3).

3. Telescopic decompositions. To develop a randomized algorithm for recovering an HSS matrix from matrix-vector products, Levitt and Martinsson [20] replaced the classical data-sparse representation (2.3) with a certain type of telescopic decomposition. A similar but different telescopic decomposition has been described in [9]. As discussed in [9] and section 3.2, this decomposition is directly linked with the data-sparse representation (2.3). In section 3.1, we introduce a general class of telescopic decompositions that includes both types. We also discuss how to convert between these types of telescopic decompositions.

Algorithm 3.1. Recovering A from a telescopic decomposition.

Require: $\{U_\tau, V_\tau, D_\tau\}_{\tau \in \mathcal{T}}$ telescopic decomposition of A for cluster tree $\mathcal{T}$ of depth L ,

Ensure: A

$A \leftarrow D_\gamma$ where γ is the root of $\mathcal{T}$

for $\ell = 1, \dots, L$ **do**

$A \leftarrow \mathbf{D}^{(\ell)} + \mathbf{U}^{(\ell)} A (\mathbf{V}^{(\ell)})^T$ $\triangleright$ with $\mathbf{D}^{(\ell)}, \mathbf{U}^{(\ell)}, \mathbf{V}^{(\ell)}$ defined as in (2.4)

end for

3.1. A general telescopic decomposition. We start with a generalization of the construction from [20], recursively defining a telescopic decomposition.

DEFINITION 3.1. Let $\mathcal{T}$ be a cluster tree of depth L for the indices $[1, \dots, n]$. A matrix $A \in \mathbb{R}^{n \times n}$ is in telescopic decomposition of telescopic rank r if there are real matrices

$$\{U_\tau, V_\tau : \tau \in \mathcal{T}, 1 \leq \text{depth}(\tau)\} \quad \text{and} \quad \{D_\tau : \tau \in \mathcal{T}\},$$

for brevity denoted by $\{U_\tau, V_\tau, D_\tau\}_{\tau \in \mathcal{T}}$ or simply $\{U_\tau, V_\tau, D_\tau\}$, with the following properties:

1. D_τ is of size $|\tau| \times |\tau|$ if $\text{depth}(\tau) = L$ and $2r \times 2r$ otherwise;
2. U_τ, V_τ have orthonormal columns; they are of size $|\tau| \times r$ if $\text{depth}(\tau) = L$ and $2r \times r$ otherwise;
3. if $L = 0$ (i.e., $\mathcal{T}$ consists only of the root γ), then $A = D_\gamma$;
4. if $L \geq 1$, then

$$(3.1) \quad A = \mathbf{D}^{(L)} + \mathbf{U}^{(L)} A^{(L-1)} (\mathbf{V}^{(L)})^T,$$

where $\mathbf{U}^{(L)}, \mathbf{V}^{(L)}, \mathbf{D}^{(L)}$ are the block diagonal matrices defined by U_τ, V_τ, D_τ as in (2.4), and the matrix

$$(3.2) \quad A^{(L-1)} := (\mathbf{U}^{(L)})^T (A - \mathbf{D}^{(L)}) \mathbf{V}^{(L)}$$

has the telescopic decomposition $\{U_\tau, V_\tau, D_\tau\}_{\tau \in \mathcal{T}_{2r}^{(L-1)}}$, where $\mathcal{T}_{2r}^{(L-1)}$ denotes a balanced cluster tree of depth $L - 1$ (see Definition 2.4).

As we will see in the following, Definition 3.1 offers significant freedom in the choice of diagonal blocks D_τ , giving rise to different types of telescopic decompositions. A simple procedure to explicitly reconstruct the matrix A from a telescopic decomposition is described in Algorithm 3.1.

3.2. From HSS matrices to telescopic decompositions. The results of this section are similar to and follow from the results in [9, section 4]. For completeness, we provide the proofs of Propositions 3.2 and 3.4 in Appendix A.

The following proposition provides a construction that turns any HSS matrix into a telescopic decomposition of the same rank.

PROPOSITION 3.2. Let

$$\{U_\tau, V_\tau, \tilde{A}_{\tau, \tau'} : \tau, \tau' \text{ sibling nodes}\} \quad \text{and} \quad \{A_{\tau, \tau} : \tau \text{ leaf node}\}$$

be the data-sparse representation defining an HSS matrix A of HSS rank r for a cluster tree $\mathcal{T}$ of depth L . For each node $\tau \in \mathcal{T}$, define

$$D_\tau := \begin{cases} A_{\tau,\tau} & \text{if } \tau \text{ is a leaf node,} \\ \begin{bmatrix} 0 & \tilde{A}_{\alpha,\beta} \\ \tilde{A}_{\beta,\alpha} & 0 \end{bmatrix} & \text{if } \tau \text{ has children } \alpha \text{ and } \beta. \end{cases}$$

Then $\{U_\tau, V_\tau, D_\tau\}$ is a telescopic decomposition of A of telescopic rank r associated with $\mathcal{T}$.

The proof of Proposition 3.2 shows that the matrix $A^{(L-1)}$ generated by the telescopic decomposition $\{U_\tau, V_\tau, D_\tau\}_{\tau \in \mathcal{T}_{2r}^{(L-1)}}$ satisfies $(A^{(L-1)})_{\tau,\tau} = D_\tau$ for every leaf τ of $\mathcal{T}_{2r}^{(L-1)}$. Letting $\mathcal{T}_{2r}^{(\ell)}$ denote a balanced cluster tree of depth $\ell \leq L-1$, we can apply this property recursively and obtain

$$(3.3) \quad (A^{(\ell)})_{\tau,\tau} = D_\tau \quad \text{for every leaf } \tau \text{ of } \mathcal{T}_{2r}^{(\ell)} \text{ for all } 1 \leq \ell \leq L-1,$$

where $A^{(\ell)}$ denotes the matrix generated by the telescopic decomposition $\{U_\tau, V_\tau, D_\tau\}_{\tau \in \mathcal{T}_{2r}^{(\ell)}}$. It follows from Proposition 3.2 and Proposition 3.4 below that telescopic decompositions with this property are in a simple one-to-one correspondence with HSS matrices, which justifies the following definition.

DEFINITION 3.3. A telescopic decomposition $\{U_\tau, V_\tau, D_\tau\}_{\tau \in \mathcal{T}}$ of a matrix A is called standard if $D_\tau = A_{\tau,\tau}$ for each τ leaf node and (3.3) holds.

PROPOSITION 3.4. Let $\mathcal{T}$ be a cluster tree of depth L and let A be the matrix generated by a standard telescopic decomposition $\{U_\tau, V_\tau, D_\tau\}_{\tau \in \mathcal{T}}$ of telescopic rank r . Then, for each nonleaf node τ with children α, β , there exist matrices $\tilde{A}_{\alpha,\beta}, \tilde{A}_{\beta,\alpha} \in \mathbb{R}^{k \times k}$ such that

$$(3.4) \quad D_\tau = \begin{bmatrix} 0 & \tilde{A}_{\alpha,\beta} \\ \tilde{A}_{\beta,\alpha} & 0 \end{bmatrix}.$$

Moreover, A is an HSS matrix of HSS rank r with the data-sparse representation

$$\{U_\tau, V_\tau : \tau \in \mathcal{T}\}, \quad \{\tilde{A}_{\tau,\tau'} : \tau, \tau' \text{ sibling nodes}\}, \quad \{A_{\alpha,\alpha} : \alpha \text{ leaf node}\}.$$

3.3. Converting a general telescopic decomposition into a standard one.

In the following, we describe a procedure that turns an arbitrary telescopic decomposition $\{U_\tau, V_\tau, D_\tau\}$ of a matrix A into a standard telescopic decomposition $\{U_\tau, V_\tau, C_\tau\}$. We stress that the matrices U_τ and V_τ remain unchanged. By the results of section 3.2, this implies the equivalence between HSS matrices of HSS rank r and matrices that admit a telescopic decomposition of telescopic rank r . In section 4, we will observe that employing the standard telescopic decomposition is preferred over other, more general telescopic decompositions because it allows us in a straightforward manner to ensure that the developed algorithm for computing Hermitian matrix functions is well defined.

For this purpose, it is crucial to understand how we can recover the principal submatrices $A_{\alpha,\alpha}$ for leaf nodes α , since these matrices correspond to the matrices C_α in the standard telescopic decomposition.

PROPOSITION 3.5. For a cluster tree $\mathcal{T}$ of depth L , let A be a matrix in telescopic decomposition $\{U_\tau, V_\tau, D_\tau\}$ of telescopic rank r . Then the following hold:

1. if $L = 0$ (i.e., $\mathcal{T}$ consists only of the root γ), then $A_{\gamma,\gamma} = D_\gamma$;
2. if $L \geq 1$, any pair of sibling leaf nodes α, β with parent τ satisfies

$$(3.5) \quad \begin{bmatrix} A_{\alpha,\alpha} & \\ & A_{\beta,\beta} \end{bmatrix} = \begin{bmatrix} D_\alpha & \\ & D_\beta \end{bmatrix} + \begin{bmatrix} U_\alpha [(A_{\tau,\tau}^{(L-1)})_{1,1}] V_\alpha^T & \\ & U_\beta [(A_{\tau,\tau}^{(L-1)})_{2,2}] V_\beta^T \end{bmatrix},$$

with the matrix $A^{(L-1)}$ from Definition 3.1, and $(A_{\tau,\tau}^{(L-1)})_{1,1}$ and $(A_{\tau,\tau}^{(L-1)})_{2,2}$ denoting the (1,1) and (2,2) diagonal blocks of $A_{\tau,\tau}^{(L-1)} \in \mathbb{R}^{2r \times 2r}$, respectively.

Proof. Point 1 follows directly from the definition of D_γ . To prove point 2, we observe that (3.1) implies

$$\begin{bmatrix} A_{\alpha,\alpha} & * \\ * & A_{\beta,\beta} \end{bmatrix} = A_{\tau,\tau} = \begin{bmatrix} D_\alpha & \\ & D_\beta \end{bmatrix} + \begin{bmatrix} U_\alpha & \\ & U_\beta \end{bmatrix} A_{\tau,\tau}^{(L-1)} \begin{bmatrix} V_\alpha^T & \\ & V_\beta^T \end{bmatrix}.$$

Therefore, taking the diagonal blocks concludes the proof. $\square$

The previous proposition combined with the fact that a telescopic decomposition of the matrix $A^{(L-1)}$ employed in (3.5) is given by $\{U_\tau, V_\tau, D_\tau\}_{\text{depth}(\tau) \leq L-1}$ (see Definition 3.1) results in a practical way to compute the matrices $A_{\alpha,\alpha}$ for all leaves α ; see Algorithm 3.2.

To satisfy the definition of a standard telescopic decomposition on the leaf level, we need to set

$$C_\alpha := A_{\alpha,\alpha}$$

for each leaf node α . Moreover, if $L \geq 1$, for each node τ of depth $L-1$ the matrix C_τ is given by $A_{\tau,\tau}^{(L-1)}$, where $A^{(L-1)}$ is now defined by (3.2) using the matrices C_α instead of D_α for a leaf α , that is,

$$\begin{aligned} A^{(L-1)} &:= (U^{(L)})^T (A - C^{(L)}) V^{(L)} = (U^{(L)})^T (A - D^{(L)}) V^{(L)} \\ &\quad + (U^{(L)})^T (D^{(L)} - C^{(L)}) V^{(L)}, \end{aligned}$$

Algorithm 3.2. Computation of principal submatrices of A given in telescopic decomposition.

Require: Matrix A in telescopic decomposition $\{U_\tau, V_\tau, D_\tau\}$ for cluster tree $\mathcal{T}$ of depth L

Ensure: $\{A_{\alpha,\alpha} : \alpha \text{ leaf node}\}$

$\hat{A}_{\gamma,\gamma} \leftarrow D_\gamma$ for root γ of $\mathcal{T}$

for $\ell = 0, \dots, L-1$ **do**

for each node τ of depth ℓ **do**

 Denoting by α, β the children of τ and defining $(\hat{A}_{\tau,\tau})_{1,1}, (\hat{A}_{\tau,\tau})_{2,2}$ as in Proposition 3.5

$$\begin{bmatrix} \hat{A}_{\alpha,\alpha} & \\ & \hat{A}_{\beta,\beta} \end{bmatrix} \leftarrow \begin{bmatrix} D_\alpha & \\ & D_\beta \end{bmatrix} + \begin{bmatrix} U_\alpha [(\hat{A}_{\tau,\tau})_{1,1}] V_\alpha^T & \\ & U_\beta [(\hat{A}_{\tau,\tau})_{2,2}] V_\beta^T \end{bmatrix}$$

end for

end for

for each leaf node α **do**

$$A_{\alpha,\alpha} \leftarrow \hat{A}_{\alpha,\alpha}$$

end for

Algorithm 3.3. Computation of a standard telescopic decomposition of A given in telescopic factors.

Require: Matrix A in telescopic decomposition $\{U_\tau, V_\tau, D_\tau\}$ for cluster tree $\mathcal{T}$ of depth L

Ensure: Standard telescopic decomposition $\{U_\tau, V_\tau, C_\tau\}$ of A
 $\{C_\alpha : \alpha \text{ leaf node}\} \leftarrow$ Algorithm 3.2 applied to $\{U_\tau, V_\tau, D_\tau\}$

for each node τ **do**

$\hat{D}_\tau \leftarrow D_\tau$

end for

for $\ell = L - 1, \dots, 0$ **do**

for each node τ of depth ℓ **do**

Denoting by α, β the children of τ

$$\hat{D}_\tau \leftarrow \hat{D}_\tau - \begin{bmatrix} U_\alpha^T (\hat{D}_\alpha - C_\alpha) V_\alpha & \\ & U_\beta^T (\hat{D}_\beta - C_\beta) V_\beta \end{bmatrix}$$

end for

$\{C_\tau : \tau \in \mathcal{T} \text{ of depth } \ell\} \leftarrow$ Algorithm 3.2 applied to $\{U_\tau, V_\tau, \hat{D}_\tau\}_{\text{depth}(\tau) \leq \ell}$

end for

with the block diagonal matrices $\mathbf{U}^{(L)}, \mathbf{V}^{(L)}, \mathbf{D}^{(L)}$, and $\mathbf{C}^{(L)}$ defined from $\{U_\tau\}$, $\{V_\tau\}$, $\{D_\tau\}$, and $\{C_\tau\}$, respectively, according to (2.4). By Definition 3.1, the matrix $(\mathbf{U}^{(L)})^T (A - \mathbf{D}^{(L)}) \mathbf{V}^{(L)}$ is generated by the telescopic decomposition $\{U_\tau, V_\tau, D_\tau\}_{\text{depth}(\tau) \leq L-1}$, hence defining

$$\hat{D}_\tau := \begin{cases} D_\tau - \begin{bmatrix} U_\alpha^T (D_\alpha - C_\alpha) V_\alpha & \\ & U_\beta^T (D_\beta - C_\beta) V_\beta \end{bmatrix} & \text{if } \tau \text{ has depth } L-1 \text{ and children } \alpha, \beta; \\ D_\tau & \text{otherwise;} \end{cases}$$

the matrix $A^{(L-1)}$ is generated by the telescopic decomposition $\{U_\tau, V_\tau, \hat{D}_\tau\}_{\text{depth}(\tau) \leq L-1}$. Therefore a standard decomposition of A can be computed by iterating Algorithm 3.2, as summarized in Algorithm 3.3. In particular, the computational complexity of transforming a telescopic decomposition into a standard one is $\mathcal{O}(r^3 2^L)$, where r is the telescopic rank of A . Assuming the threshold size and the telescopic rank to be constant, this shows linear complexity in the size of A .

3.4. Symmetric telescopic decompositions. For a symmetric matrix, the definition of a telescopic decomposition can be adjusted to reflect symmetry.

DEFINITION 3.6. A telescopic decomposition $\{U_\tau, V_\tau, D_\tau\}_{\tau \in \mathcal{T}}$ is said to be symmetric if $V_\tau = U_\tau$ and $D_\tau^T = D_\tau$ hold for every $\tau \in \mathcal{T}$. In analogy to the nonsymmetric case, we employ the term standard if (3.3) is satisfied. For simplicity, a symmetric telescopic decomposition is denoted by $\{U_\tau, D_\tau\}_{\tau \in \mathcal{T}}$, avoiding the repetition of U_τ .

If A is symmetric and has HSS rank r , there exist data-sparse representations of the form (2.3), for which $U_\tau = V_\tau$ have r columns and $\tilde{A}_{\tau, \tau'}^T = \tilde{A}_{\tau', \tau}$ holds for each pair of sibling nodes τ, τ' (see [21, section 4.1]). Therefore, Proposition 3.2 implies that A admits a symmetric telescopic decomposition of rank r . We also recall that the procedure described in section 3.3 converts a telescopic decomposition $\{U_\tau, V_\tau, D_\tau\}$ into a standard one, without changing the matrices U_τ, V_τ . In turn, the same procedure can be employed to convert a symmetric telescopic decomposition into a standard symmetric telescopic decomposition.

4. Computing telescopic decompositions for functions of symmetric HSS matrices. If A is a symmetric HSS matrix with spectrum contained in $[\lambda_{\min}, \lambda_{\max}]$ and the function f is analytic on $[\lambda_{\min}, \lambda_{\max}]$, then $f(A)$ can usually be well approximated by an HSS matrix. While this has been observed before [6, section 3], it is nontrivial to develop an algorithm that fully exploits this property. In this section, we derive such an algorithm that computes a telescopic decomposition for an HSS approximation of $f(A)$ starting from a standard symmetric telescopic decomposition of A . If needed, this can be converted into a standard telescopic decomposition, employing the results of section 3.3, and therefore into an HSS data-sparse representation (2.3). If f is a rational function of a certain degree, we show that such an approximation is exact and, otherwise, the approximation error is bounded using a rational approximation of f on $[\lambda_{\min}, \lambda_{\max}]$.

Block rational Krylov subspaces play a key role in our construction, and we briefly recall their definition; see [7] for more details.

DEFINITION 4.1. Let $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times b}$, and let $\xi_k = (\xi_0, \dots, \xi_{k-1})$, with $\xi_j \in \mathbb{C} \cup \{\infty\}$, be a list of poles that are not eigenvalues of A . Then the associated rational Krylov subspace is defined as

$$\mathcal{Q}(A, B, \xi_k) := \left\{ q(A)^{-1} \sum_{j=0}^{k-1} A^j B C_j, \quad \text{with } C_j \in \mathbb{R}^{b \times b} \right\},$$

where

$$q(x) := \prod_{\xi \in \xi_k, \xi \neq \infty} (x - \xi).$$

An orthonormal basis of a rational Krylov subspace can be computed by the *rational Arnoldi method* [7, Algorithm 2.1]. During the procedure, shifted linear systems with the matrices² $A - \xi_j I$ need to be solved, and this can become computationally expensive for a large-scale matrix A . One novelty of this work is to employ rational Krylov subspaces associated with small-size matrices for the computation of large-scale matrix functions. Proposition 4.2 below is a simple but key result that allows us to do this. We recall (2.4): Given a family of matrices $(C_i \in \mathbb{R}^{m_i \times n_i})$, the matrix $\mathbf{C} = \text{blkdiag}(C_i)$ is the block diagonal matrix containing the matrices C_i as diagonal blocks (in their natural order).

PROPOSITION 4.2. Let $\mathbf{A} := \text{blkdiag}(A_i)$ and $\mathbf{B} := \text{blkdiag}(B_i)$ with $A_i \in \mathbb{R}^{n_i \times n_i}$ and $B_i \in \mathbb{R}^{n_i \times b_i}$ and let ξ_k be a list of poles. Then an orthonormal basis of $\mathcal{Q}(\mathbf{A}, \mathbf{B}, \xi_k)$ is given by $\mathbf{U} := \text{blkdiag}(U_i)$, where U_i is an orthonormal basis of $\mathcal{Q}(A_i, B_i, \xi_k)$.

Proof. By Definition 4.1, every element of $\mathcal{Q}(\mathbf{A}, \mathbf{B}, \xi_k)$ takes the form

$$q(\mathbf{A})^{-1} \sum_{j=0}^{k-1} \mathbf{A}^j \mathbf{B} C_j = \sum_{j=0}^{k-1} \text{blkdiag}(q(A_i)^{-1} A_i^j B_i) C_j$$

for some matrix $C_j \in \mathbb{R}^{b \times b}$ with $b = \sum_i b_i$. Because $q(A_i)^{-1} A_i^j B_i \in \mathcal{Q}(A_i, B_i, \xi_k)$ for every i, j , it follows that $\mathcal{Q}(\mathbf{A}, \mathbf{B}, \xi_k) = \text{span}(\text{blkdiag}(U_i))$. $\square$

Summarizing results from [1], the following theorem shows how low-rank updates of a matrix function can be approximated using rational Krylov subspaces.

²The linear system is replaced by a matrix product with A if $\xi_j = \infty$.

THEOREM 4.3 (see [1, Theorem 4.5]). Let $A = C + Z\tilde{A}Z^T$, where $C \in \mathbb{R}^{n \times n}$, $\tilde{A} \in \mathbb{R}^{m \times m}$ are symmetric and $Z \in \mathbb{R}^{n \times m}$. Let $[\lambda_{\min}, \lambda_{\max}]$ be an interval that contains the spectra of A and C , and let $\xi_k = (\xi_0, \dots, \xi_{k-1})$, with $\xi_j \in \mathbb{C} \cup \{\infty\}$, be a list of poles closed under complex conjugation. For a function f analytic on $[\lambda_{\min}, \lambda_{\max}]$, we consider the approximation

$$(4.1) \quad f(A) \approx \hat{F} := f(C) + W(f(W^T A W) - f(W^T C W))W^T,$$

where W is an orthonormal basis of $\mathcal{Q}(C, Z, \xi_k)$. Then the approximation error $E(f) := \hat{F} - f(A)$ satisfies

$$\|E(f)\|_2 \leq 4 \min_{r \in \mathbb{P}_k/q_k} \|f - r\|_\infty,$$

where $q_k = \prod_{\xi \in \xi_k, \xi \neq \infty} (x - \xi)$, $\mathbb{P}_k$ denotes the set of polynomials of degree at most $k - 1$, and $\|\cdot\|_\infty$ denotes the supremum norm on $[\lambda_{\min}, \lambda_{\max}]$. In particular, the approximation (4.1) is exact if $f \in \mathbb{P}_k/q_k$.

We will apply Theorem 4.3 recursively to telescopic decompositions. In order to do so conveniently, we slightly loosen our assumptions on a standard telescopic decomposition. Given a cluster tree $\mathcal{T}$ associated with $[1, \dots, n]$ we assume that a matrix $A \in \mathbb{R}^{n \times n}$ can be written as

$$(4.2) \quad A = \tilde{C}^{(L)} + Z^{(L)} A^{(L-1)} (Z^{(L)})^T,$$

where

- $A^{(L-1)}$ admits a standard symmetric telescopic decomposition $\{U_\tau, C_\tau\}_{\tau \in \mathcal{T}_{2r}^{(L-1)}}$;
- $\tilde{C}^{(L)} = \text{blkdiag}(\tilde{C}_\tau : \tau \text{ leafnode in } \mathcal{T})$, and $\tilde{C}_\tau = A_{\tau\tau}$;
- $Z^{(L)} = \text{blkdiag}(Z_\tau : \tau \text{ leafnode in } \mathcal{T})$, with $Z_\tau \in \mathbb{R}^{|\tau| \times r}$.

The key difference to assuming that A has a standard symmetric telescopic decomposition is that no orthogonality is enforced on $Z^{(L)}$, the factors on the leaf level. The key advantage of (4.2) is that it remains unaffected when multiplying with certain block diagonal matrices. The following proposition additionally shows how to move one level up.

PROPOSITION 4.4. Let A be an $n \times n$ matrix admitting the decomposition (4.2), and $W^{(L)} = \text{blkdiag}(W_\tau : \tau \text{ leaf node})$ with $W_\tau \in \mathbb{R}^{|\tau| \times m}$. Then,

$$(W^{(L)})^T A W^{(L)} = \tilde{C}^{(L-1)} + Z^{(L-1)} A^{(L-2)} (Z^{(L-1)})^T,$$

with the block diagonal matrices

$$\tilde{C}^{(L-1)} = \text{blkdiag}(\tilde{C}_\tau : \text{depth}(\tau) = L - 1), \quad Z^{(L-1)} = \text{blkdiag}(Z_\tau : \text{depth}(\tau) = L - 1)$$

containing the diagonal blocks

$$Z_\tau = \begin{bmatrix} W_\alpha^T Z_\beta & \\ & W_\alpha^T Z_\beta \end{bmatrix} U_\tau \in \mathbb{R}^{2m \times r},$$

$$\tilde{C}_\tau = \begin{bmatrix} W_\alpha^T \tilde{C}_\alpha W_\alpha & & \\ & W_\beta^T \tilde{C}_\beta W_\beta \end{bmatrix} + \begin{bmatrix} W_\alpha^T Z_\alpha & & \\ & W_\beta^T Z_\beta \end{bmatrix} C_\tau \begin{bmatrix} W_\alpha^T Z_\alpha & & \\ & W_\beta^T Z_\beta \end{bmatrix}^T$$

$$\in \mathbb{R}^{2m \times 2m},$$

and the matrix $A^{(L-2)}$ generated by the standard telescopic decomposition $\{U_\tau, C_\tau\}_{\tau \in \mathcal{T}_{2r}^{(L-2)}}$, where the matrices U_τ and C_τ stem from the telescopic decomposition of $A^{(L-1)}$. Moreover, $\tilde{C}_\tau = A_{\tau,\tau}$ holds for each leaf node $\tau \in \mathcal{T}_{2m}^{(L-1)}$.

Proof. Considering (4.2) and applying the telescopic decomposition (3.1) of the matrix $A^{(L-1)}$, we get

$$\begin{aligned} & (\mathbf{W}^{(L)})^T A \mathbf{W}^{(L)} \\ &= (\mathbf{W}^{(L)})^T \left(\tilde{\mathbf{C}}^{(L)} + \mathbf{Z}^{(L)} (\mathbf{C}^{(L-1)} + \mathbf{U}^{(L-1)} A^{(L-2)} (\mathbf{U}^{(L-1)})^T) (\mathbf{Z}^{(L)})^T \right) \mathbf{W}^{(L)} \\ &= \tilde{\mathbf{C}}^{(L-1)} + \mathbf{Z}^{(L-1)} A^{(L-2)} (\mathbf{Z}^{(L-1)})^T \end{aligned}$$

with $\mathbf{C}^{(L-1)}, \mathbf{U}^{(L-1)}$ defined from $\{C_\tau\}, \{U_\tau\}$, in accordance with (2.4). Moreover, because $\{U_\tau, C_\tau\}_{\tau \in \mathcal{T}_{2r}^{(L-1)}}$ is a standard decomposition of $A^{(L-1)}$, the relation $\tilde{\mathbf{C}}_\tau = A_{\tau, \tau}$ holds for each leaf node $\tau \in \mathcal{T}_{2m}^{(L-1)}$. $\square$

To approximate $f(A)$ for a matrix A admitting the decomposition (4.2), we use the construction of Theorem 4.3 with $C = \tilde{\mathbf{C}}^{(L)}$ and $Z = \mathbf{Z}^{(L)}$, which yields the approximation

$$(4.3) \quad f(A) \approx f(\tilde{\mathbf{C}}^{(L)}) + \mathbf{W}^{(L)} \left[f((\mathbf{W}^{(L)})^T A \mathbf{W}^{(L)}) - f((\mathbf{W}^{(L)})^T \tilde{\mathbf{C}}^{(L)} \mathbf{W}^{(L)}) \right] (\mathbf{W}^{(L)})^T,$$

where $\mathbf{W}^{(L)}$ is an orthogonal basis for $\mathcal{Q}(\tilde{\mathbf{C}}^{(L)}, \mathbf{Z}^{(L)}, \boldsymbol{\xi}_k)$. Note that the three evaluations of f are all well defined³ because $\tilde{\mathbf{C}}^{(L)}$ contains diagonal blocks of A and $(\mathbf{W}^{(L)})^T \tilde{\mathbf{C}}^{(L)} \mathbf{W}^{(L)}$, $(\mathbf{W}^{(L)})^T A \mathbf{W}^{(L)}$ are orthogonal compressions. By eigenvalue interlacing, the spectra of these three matrices are contained in $[\lambda_{\min}, \lambda_{\max}]$. We now make two observations:

(i) Proposition 4.2 implies that the $n \times 2^L r k$ matrix $\mathbf{W}^{(L)}$ takes the form

$\mathbf{W}^{(L)} = \text{blkdiag}(W_\tau : \tau \text{ leaf of } \mathcal{T}),$ with $W_\tau \in \mathbb{R}^{|\tau| \times r k}$ orthonormal basis of $\mathcal{Q}(\tilde{\mathbf{C}}_\tau, Z_\tau, \boldsymbol{\xi}_k)$.

(ii) Proposition 4.4 implies that $B^{(L-1)} := (\mathbf{W}^{(L)})^T A \mathbf{W}^{(L)}$ admits the decomposition

$$(4.4) \quad B^{(L-1)} = \tilde{\mathbf{C}}^{(L-1)} + \mathbf{Z}^{(L-1)} A^{(L-2)} (\mathbf{Z}^{(L-1)})^T.$$

In (4.3), the function f needs to be evaluated for three matrices. This is cheap for $\tilde{\mathbf{C}}^{(L)}$ and $\mathbf{W}^{(L)} \tilde{\mathbf{C}}^{(L)} \mathbf{W}^{(L)}$ because these matrices are block diagonal with small diagonal blocks, for which the evaluation of f is computed explicitly. The expensive part is the evaluation of f for $B^{(L-1)} = (\mathbf{W}^{(L)})^T A \mathbf{W}^{(L)}$. For this purpose, we use the decomposition (4.4) and apply the approximation (4.3) again:

$$\begin{aligned} f(B^{(L-1)}) &\approx f(\tilde{\mathbf{C}}^{(L-1)}) \\ &+ \mathbf{W}^{(L-1)} \left[f(B^{(L-2)}) - f((\mathbf{W}^{(L-1)})^T \tilde{\mathbf{C}}^{(L-1)} \mathbf{W}^{(L-1)}) \right] (\mathbf{W}^{(L-1)})^T, \end{aligned}$$

where $B^{(L-2)} := (\mathbf{W}^{(L-1)})^T B^{(L-1)} \mathbf{W}^{(L-1)}$ and

$$\begin{aligned} & \mathbf{W}^{(L-1)} \\ &= \text{blkdiag}(W_\tau : \tau \text{ leaf of } \mathcal{T}_{2rk}^{(L-1)}), \text{ with } W_\tau \text{ orthonormal basis of } \mathcal{Q}(\tilde{\mathbf{C}}_\tau, Z_\tau, \boldsymbol{\xi}_k). \end{aligned}$$

³This is the main reason for using a standard telescopic decomposition of A . Indeed, for a nonstandard telescopic decomposition, the well-defined evaluation of f is not guaranteed.

Since $\tilde{\mathbf{C}}^{(L-1)}$ contains diagonal blocks of $B^{(L-1)}$, the three evaluations of f are again well defined.

The procedure described above is repeated recursively until reaching a tree of depth 0, and at that point one simply computes the matrix function of the corresponding dense matrix of size $2rk \times 2rk$ explicitly. The i th step of the recursive procedure proceeds as follows: one assumes a decomposition of the form

$$(4.5) \quad B^{(L-i)} = \tilde{\mathbf{C}}^{(L-i)} + \mathbf{Z}^{(L-i)} A^{(L-i-1)} (\mathbf{Z}^{(L-i)})^T,$$

where $\tilde{\mathbf{C}}^{(L-i)}$, $\mathbf{Z}^{(L-i)}$ are block diagonal matrices defined from $\{\tilde{C}_\tau\}$, $\{Z_\tau\}$, in accordance with (2.4), and $A^{(L-i-1)}$ admits a standard symmetric telescopic decomposition $\{U_\tau, C_\tau\}_{\tau \in \mathcal{T}_{2r}^{(L-i-1)}}$. Then $f(B^{(L-i)})$ is approximated by

$$(4.6) \quad f(\tilde{\mathbf{C}}^{(L-i)}) + \mathbf{W}^{(L-i)} \left[f(B^{(L-i-1)}) - f((\mathbf{W}^{(L-i)})^T \tilde{\mathbf{C}}^{(L-i)} \mathbf{W}^{(L-i)}) \right] (\mathbf{W}^{(L-i)})^T,$$

where $B^{(L-i-1)} := (\mathbf{W}^{(L-i)})^T B^{(L-i)} \mathbf{W}^{(L-i)}$ and

$$\mathbf{W}^{(L-i)}$$

$$= \text{blkdiag}(W_\tau : \tau \text{ leaf of } \mathcal{T}_{2rk}^{(L-i)}), \text{ with } W_\tau \text{ orthonormal basis of } \mathcal{Q}(\tilde{C}_\tau, Z_\tau, \xi_k),$$

explicitly computing $f(\tilde{\mathbf{C}}^{(L-i)})$, $f((\mathbf{W}^{(L-i)})^T \tilde{\mathbf{C}}^{(L-i)} \mathbf{W}^{(L-i)})$ and recursively approximating $f(B^{(L-i-1)})$. In particular, the procedure can be iterated considering the decomposition

$$B^{(L-i-1)} = \tilde{\mathbf{C}}^{(L-i-1)} + \mathbf{Z}^{(L-i-1)} A^{(L-i-2)} (\mathbf{Z}^{(L-i-1)})^T,$$

given by Proposition 4.4.

Assuming that A admits a standard (symmetric) telescopic decomposition $\{U_\tau, C_\tau\}$, the described procedure starts by taking $Z_\tau = U_\tau$ and $\tilde{C}_\tau = C_\tau$ for each leaf node τ . It results in a telescopic decomposition for an approximation of $f(A)$, with generators $\{W_\tau, D_\tau\}$; the generators W_τ are defined by the rational Krylov subspaces constructed throughout the process, whereas D_τ takes the form

$$D_\tau := \begin{cases} f(\tilde{C}_\tau) & \text{if } \tau \text{ is the root node,} \\ f(\tilde{C}_\tau) - W_\tau f(W_\tau^T \tilde{C}_\tau W_\tau) W_\tau^T & \text{otherwise.} \end{cases}$$

The procedure is summarized in Algorithm 4.1.

Let us emphasize that even if the decomposition (4.2) is used in the course of the algorithm, the final result has a symmetric telescopic decomposition in the sense of Definition 3.6. The nonorthogonal factors $\mathbf{Z}^{(\ell)}$ are only needed to represent intermediate stages.

To discuss the complexity of Algorithm 4.1, let r denote the HSS/telescopic rank of A , and assume that $n = 2^L t$, with the threshold size $t \sim kr$ and that the cluster tree is balanced: $\mathcal{T} = \mathcal{T}_t^{(L)}$. On the leaf level L , the computation of all W_τ and C_τ requires $\mathcal{O}(2^L(t^3 + krt^2)) = \mathcal{O}(nk^2r^2)$ operations. On level $\ell < L$, the complexity is $\mathcal{O}(2^\ell k^3 r^3) = \mathcal{O}(2^{-(L-\ell)} nk^2 r^2)$. This gives a total complexity of

$$(4.7) \quad \mathcal{O}(nk^2r^2),$$

which is linear in n if both r and k are considered constant.

We conclude this section with a result that bounds the approximation error of Algorithm 4.1 by the rational approximation error of f on $[\lambda_{\min}, \lambda_{\max}]$.

Algorithm 4.1. Computation of $f(A)$ for symmetric HSS matrix A in telescopic decomposition.

Require: $\{U_\tau, C_\tau\}$ standard symmetric telescopic decomposition of matrix A , function f , list of poles $\xi_k = (\xi_0, \dots, \xi_{k-1}) \subseteq \mathbb{C} \cup \{\infty\}$ closed under complex conjugation

Ensure: $\{W_\tau, D_\tau\}$ telescopic factorization of an approximation to $f(A)$

```

for  $\ell = L, L-1, \dots, 0$  do
  for each node  $\tau$  of depth  $\ell$  do
    if  $\ell = L$  then
       $Z_\tau \leftarrow U_\tau$ 
       $\tilde{C}_\tau \leftarrow C_\tau$ 
    else
      Let  $\alpha$  and  $\beta$  be the children of  $\tau$ 
       $Z_\tau \leftarrow \begin{bmatrix} W_\alpha^T Z_\alpha & W_\beta^T Z_\beta \end{bmatrix} U_\tau$ 
       $\tilde{C}_\tau \leftarrow \begin{bmatrix} W_\alpha^T \tilde{C}_\alpha W_\alpha & W_\beta^T \tilde{C}_\beta W_\beta \end{bmatrix} + \begin{bmatrix} W_\alpha^T Z_\alpha & W_\beta^T Z_\beta \end{bmatrix} C_\tau \begin{bmatrix} Z_\alpha^T W_\alpha & Z_\beta^T W_\beta \end{bmatrix}$ 
    end if
    if  $\ell = 0$  then
       $D_\tau \leftarrow f(\tilde{C}_\tau)$ 
    else
       $W_\tau \leftarrow$  orthonormal basis of  $\mathcal{Q}(\tilde{C}_\tau, Z_\tau, \xi_k)$ 
       $D_\tau \leftarrow f(\tilde{C}_\tau) - W_\tau f(W_\tau^T \tilde{C}_\tau W_\tau) W_\tau^T$ 
    end if
  end for
end for

```

THEOREM 4.5. *With the notation introduced in Theorem 4.3, let A be a symmetric HSS matrix associated with a cluster tree $\mathcal{T}$ of depth L . Let f be analytic on an interval $[\lambda_{\min}, \lambda_{\max}]$ containing the eigenvalues of A . Letting $E(f)$ denote the difference between $f(A)$ and the output of Algorithm 4.1 applied to a standard decomposition of A , it holds that*

$$(4.8) \quad \|E(f)\|_2 \leq 4L \min_{r \in \mathbb{P}_k/q_k} \|f - r\|_\infty.$$

Proof. Let $\{W_\tau, D_\tau\}_{\tau \in \mathcal{T}}$ be the telescopic decomposition returned by Algorithm 4.1 applied to a standard symmetric telescopic decomposition $\{U_\tau, C_\tau\}_{\tau \in \mathcal{T}}$. For each $1 \leq i \leq L-1$ let $B^{(L-i)}$ be the matrices defined in (4.5), and for each leaf node $\tau \in \mathcal{T}_{2rk}^{(L-i)}$, let $\tilde{C}_\tau = B_{\tau,\tau}^{(L-i)}$. Moreover, to streamline the notation, we let $B^{(L)} := A$ and $\tilde{C}_\tau := C_\tau$ for each leaf node $\tau \in \mathcal{T}$. For each i , let $E^{(L-i)}(f)$ be the difference between $f(B^{(L-i)})$ and its approximation (4.6). Since $[\lambda_{\min}, \lambda_{\max}]$ contains the eigenvalues of $B^{(L-i)}$ and $\tilde{C}_\tau$ for each τ , Theorem 4.3 implies

$$(4.9) \quad \|E^{(L-i)}(f)\|_2 \leq 4 \min_{r \in \mathbb{P}_k/q_k} \|f - r\|_\infty.$$

Denoting by $F^{(L-i)}$ the matrix generated by the symmetric telescopic decomposition $\{W_\tau, D_\tau\}_{\tau \in \mathcal{T}_{2rk}^{L-i}}$ for $0 \leq i \leq L$, we have

$$f(B^{(L-i)}) - F^{(L-i)} = \begin{cases} 0 & \text{if } i = L; \\ f(B^{(L-i)}) - D^{(L-i)} - W^{(L-i)} F^{(L-i-1)} (W^{(L-i-1)})^T & \text{otherwise,} \end{cases}$$

where $D^{(L-i)}$, $W^{(L-i)}$ are the block diagonal matrices defined by $\{D_\tau\}$, $\{W_\tau\}$, in accordance with (2.4). For $i < L$ we observe that

$$f(B^{(L-i)}) - D^{(L-i)} = W^{(L-i)} f(B^{(L-i-1)}) (W^{(L-i)})^T + E^{(L-i)}(f),$$

and, hence,

$$f(B^{(L-i)}) - F^{(L-i)} = W^{(L-i)} (f(B^{(L-i-1)}) - F^{(L-i-1)}) (W^{(L-i)})^T + E^{(L-i)}(f).$$

Using that the matrices $W^{(L-i)}$ have orthonormal columns, this implies

$$\|E(f)\|_2 = \|f(B^{(L)}) - F^{(L)}\|_2 \leq \sum_{i=0}^{L-1} \|E^{(L-i)}(f)\|_2,$$

which concludes the proof after applying the inequality (4.9). $\square$

The upper bound (4.8) of Theorem 4.5 usually overestimates the error, as the interval $[\lambda_{\min}, \lambda_{\max}]$ represents a relaxation of the spectra of the matrices involved in the construction. Although difficult to quantify rigorously, preliminary numerical experiments indicate that Algorithm 4.1 benefits from a phenomenon sometimes called spectral adaptation [13] and performs significantly better than predicted by the bound in situations where, for example, the matrix A has a few eigenvalues close to a singularity of f .

5. Pole selection. Theorem 4.5 shows that the choice of poles in the rational Krylov subspaces $\mathcal{Q}(\tilde{C}_\tau, Z_\tau, \xi_k)$ is critical to the convergence of Algorithm 4.1. Normally, repeated poles are preferred to reduce the cost of solving the shifted linear systems needed for constructing a basis of the subspace [13]. However, such considerations do not apply to Algorithm 4.1; solving linear systems with the small matrix $\tilde{C}_\tau$ (shifted by a pole) is cheap.

By Theorem 4.5, if f is an analytic function on an interval $[a, b]$ that contains the eigenvalues of A and the list of poles ξ_k satisfies

$$(5.1) \quad \min_{r \in \mathbb{P}_k/q_k} \|f - r\|_\infty \leq \frac{\epsilon}{4L},$$

then Algorithm 4.1 returns an approximation of $f(A)$ within an error bounded by a user-specified tolerance $\epsilon > 0$. In the following, we describe explicit pole selection strategies that ensure (5.1) for two important classes of functions. For general f , general rational approximation methods, like the AAA algorithm [25], can be used to select the poles.

5.1. Exponential function. In the context of the matrix exponential, it is not uncommon to use polynomial approximations, that is, all poles are infinite. However, the corresponding (polynomial) Krylov subspace methods often converge poorly when the spectrum is wide, that is, $a \ll b$; see [2, 18] for theoretical results. As their computational overhead is small in our setting, it is preferable to use rational approximations/Krylov subspaces. Assuming $b \leq 0$ (which can always be attained by

shifting the matrix), it is well known [10] that for every k there exist a list of poles ξ_k and a (universal) constant C such that

$$\min_{r \in \mathbb{P}_k/q_k} \|f - r\|_\infty \leq CK_e^{-k}, \quad K_e \approx 9.289.$$

In turn, this means that it suffices to choose

$$k \geq \log(4LC\epsilon^{-1})/\log(K_e)$$

such that Algorithm 4.1 applied to a symmetric negative semidefinite HSS matrix A with nonpositive eigenvalues returns an approximation with an error ϵ . In particular, note that these estimates are independent of the width of the spectrum. Following (4.7), this gives a complexity of

$$\mathcal{O}(n(\log \log n + \log \epsilon^{-1})^2 r^2),$$

where r is the telescopic/HSS rank of A . For example, this implies that the fixed-accuracy approximation to the exponential of *any* tridiagonal symmetric negative semidefinite matrix A has nearly linear complexity $\mathcal{O}(n(\log \log n)^2)$. We are not aware of any other algorithm that can achieve this.

5.2. Markov functions. Following the exposition in [1], we discuss pole selection for Markov functions, i.e., functions that can be represented as

$$(5.2) \quad f(z) = \int_\alpha^\beta \frac{d\mu(x)}{z - x}$$

for some positive measure $\mu(x)$ and $-\infty \leq \alpha < \beta < \infty$. Important examples of functions in this class are

$$\frac{\log(1+z)}{z} = \int_{-\infty}^{-1} \frac{-1/x}{z-x} dx \quad \text{and} \quad z^\gamma = \frac{\sin(\pi\gamma)}{\pi} \int_{-\infty}^0 \frac{|x|^\gamma}{z-x} dx,$$

with $-1 < \gamma < 0$.

Now, let f be a Markov function (5.2) with $\beta \leq 0$ and let A be a symmetric positive definite HSS matrix, with the smallest and largest eigenvalues given by $a > 0$ and $b \geq a$, respectively. The quasi-optimal rational approximation of f has been discussed in, e.g., [2, section 6.2], which for every k provides a list of poles ξ_k such that

$$\min_{r \in \mathbb{P}_k/q_k} \|f - r\|_\infty \leq 4\|f\|_\infty \exp\left(-k \frac{\pi^2}{\log(16b/a)}\right).$$

Hence, to achieve accuracy ϵ in Algorithm 4.1, one can choose

$$(5.3) \quad k \geq \log(16L\|f\|_\infty\epsilon^{-1}) \log(16b/a)/\pi^2.$$

Assuming a polynomial growth of the condition number b/a with respect to n , the complexity of the algorithm for Markov functions is, according to (4.7), given by

$$\mathcal{O}(n \log^2 n (\log \log n + \log \epsilon^{-1} + \log \|f\|_\infty)^2 r^2).$$

For example, the fixed-accuracy approximation to the inverse square root $A^{-1/2}$ requires nearly $\mathcal{O}(n \log^4 n)$ operations, when ignoring the $\log \log n$ factor and assuming the HSS rank r of A to be constant.

6. Numerical experiments. We have implemented Algorithm 4.1 in MATLAB and have made the code freely accessible at <https://github.com/numpi/HSS-matfun>; this implementation will be denoted by **TelFun** in the following. In our implementation, we allow for variable HSS/telescopic ranks (see Remark 2.2) and employ deflation criteria in the computation of orthonormal bases for rational Krylov subspaces, removing vectors that after the (re)orthogonalization step have a norm smaller than a prescribed tolerance, proportional to the required accuracy. The threshold size t of the employed HSS matrices is fixed to 256. In our experiments, all standard operations with HSS matrices, such as matrix-vector products, have been performed using the **hm-toolbox** [23]. In the tables presented in this section, columns with the caption “err” denote the relative error in the Frobenius norm, compared with the result computed by a standard dense solver. Columns with the caption “time” report the observed execution time in seconds. All experiments have been executed on a server with two Intel(R) Xeon(R) E5-2650v4 CPU running at 2.20 GHz and 256 GB of RAM, using MATLAB R2021a with the Intel(R) Math Kernel Library Version 2019.0.3.

The main competitor, denoted by **CKM**, is the algorithm developed in [6], in which the authors use the HSS structure of A to perform a divide-and-conquer method for the computation of $f(A)$. The algorithm computes rational Krylov subspaces associated with (possibly large) HSS matrices, exploiting the structure in solving linear systems. The algorithm can also monitor the variation of the norm of the solution when a new pole is employed; this quantity can be used to stop the procedure if the desired accuracy is reached. We utilized an implementation of this algorithm available at <https://github.com/Alice94/MatrixFunctions-Banded-HSS>.

6.1. Computation of the inverse. Algorithm 4.1 is an attractive method for computing the inverse of a symmetric positive definite HSS matrix. By Theorem 4.5, this algorithm returns the exact inverse (at least in exact arithmetic) when employing only one zero pole. It is worth mentioning that in this case, the update formula (4.1) reduces to the Sherman–Morrison formula; see also [1, section 3.3]. In turn, it follows that Algorithm 4.1 is equivalent to the algorithm described in [9]. We have tested **TelFun** in this situation for two different matrices. In Figure 2, we report the results for the inversion of the discretized Laplacian, that is,

$$(6.1) \quad A = -\frac{1}{h^2} \begin{bmatrix} -2 & 1 & & & \\ 1 & \ddots & \ddots & & \\ & \ddots & \ddots & 1 & \\ & & & 1 & -2 \end{bmatrix} \in \mathbb{R}^{n \times n},$$

where $h = \frac{1}{n+1}$. In Table 1, we show the results obtained when inverting a more general HSS matrix (which is not banded) given by the Grünwald–Letnikov finite difference discretization of the symmetric fractional derivative operator $L_\alpha := \frac{\partial^\alpha}{\partial x^\alpha}$ [24] for $\alpha = 1.5$. In contrast to (6.1), this finite difference approximation does not yield a sparse matrix, coherently with the nonlocal properties of fractional differential operators. It can be proven that the matrix can be approximated in the HSS format [22] with an HSS rank $\mathcal{O}(\log n)$.

Additionally to **CKM**, we also compare to the randomized algorithm introduced by Levitt and Martinsson in [20] (denoted by **LM**) based on the solution of a small number of linear systems involving A and the **inv** procedure for HSS matrices implemented in the **hm-toolbox** which is based on the ULV factorization described in [5] and explained in [23, section 4.3]. The HSS ranks are calculated using the **hssrank** command from [23], employing the default tolerance of 10^{-12} .




FIG. 2. Comparison of the newly proposed algorithm `TelFun` with CKM, LM, and the `inv` command of the `hm-toolbox` based on the ULV decomposition, for computing A^{-1} , where A is the discretized Laplacian (6.1).

TABLE 1

Comparison of the newly proposed algorithm `TelFun` with CKM, LM, and the `inv` command of the `hm-toolbox` based on the ULV decomposition, for computing A^{-1} , where A is the Grünwald-Letnikov finite difference discretization of the fractional derivative of order $\alpha = 1.5$ [22, 24].

n	HSS rank A	time <code>TelFun</code>	time CKM	time LM	time ULV	time Dense
1024	29	0.11	0.65	0.31	0.26	0.08
2048	32	0.21	0.84	0.44	0.38	0.37
4096	35	0.36	2.17	1.06	0.63	1.72
8192	37	0.62	5.48	2.14	1.41	11.20

n	HSS rank A	err <code>TelFun</code>	err CKM	err LM	err ULV
1024	29	5.50e-13	1.11e-12	4.16e-13	7.24e-12
2048	32	4.78e-13	3.26e-12	1.18e-12	2.11e-11
4096	35	2.08e-12	2.75e-11	2.53e-12	6.81e-11
8192	37	4.99e-12	4.36e-11	9.68e-12	1.75e-10

Although designed for general matrix functions, the results demonstrate that `TelFun` is also competitive for the special case of matrix inversion, in terms of speed and accuracy. Even the closest competitor ULV is significantly slower, by up to a factor 3–4. We have verified that the (larger) error observed for ULV is entirely due to the effects of roundoff error.

6.2. Computation of the exponential function. To show the effectiveness of rational approximation of the exponential function, in Table 2 we compute the matrix exponential of a tridiagonal matrix A , whose eigenvalues are uniformly dis-

tributed in $[-10^a, 0]$, for different values of a . For the computation, we compare the presented method with optimal poles and CKM using rational Krylov subspaces (with optimal poles) and/or polynomial Krylov subspaces; in the latter case the built-in stopping criteria are employed. In Figure 3 we also report the comparison between the presented method and the `expm` function implemented in the `hm-toolbox` for the computation of $\exp(A)$ based on the Padé approximant, where A is the discretized Laplacian defined in (6.1).

Again, our newly proposed method **TelFun** is significantly faster than the competitors, while resulting in comparable accuracy. Note that CKM Poly appears to not use the correct stopping criterion for larger a , resulting in an unacceptably large error. The observation that CKM Rat is slower than **TelFun** can be explained by the fact that CKM passes through the HODLR format, which incurs higher computational complexity. The Padé approximant employed in `expm` is evaluated using the ULV

TABLE 2
Computation of the matrix exponential of a matrix of size 4096, whose eigenvalues are uniformly distributed in $[-10^a, 0]$, for different values of a . The accuracy is set to 10^{-8} .

a	time TelFun	time CKM Poly	time CKM Rat	err TelFun	err CKM Poly	err CKM Rat
0	0.82	1.38	13.06	1.04e-11	2.16e-10	1.52e-10
2	0.67	1.31	11.15	2.89e-10	1.75e-07	6.11e-09
4	0.76	0.48	10.32	2.50e-12	1.05e-03	1.20e-08
6	0.52	0.52	10.13	2.84e-10	1.03e-02	4.69e-11
8	0.53	0.48	10.14	3.36e-08	1.21e + 01	5.28e-08




FIG. 3. Computation of $\exp(A)$, where A is the discretized Laplacian (6.1) using **TelFun** and the routine `expm` of the `hm-toolbox`. The accuracy is set to 10^{-8} .

TABLE 3

Comparison of the newly proposed algorithm **TelFun** (using optimal poles), **CKM** with extended Krylov subspaces, and the evaluation of a rational approximation, for the computation of $f(A)$ with accuracy of 10^{-8} , where $f(x) = 1/\sqrt{x}$, and A is the sampling from a Gaussian Markov random field.

n	HSS rank A	time TelFun	time CKM	time Rat	time Dense
512	22	0.13	0.22	1.28	0.03
1024	20	0.29	0.30	2.38	0.23
2048	21	0.57	0.73	5.09	0.87
4096	21	1.24	1.55	12.82	7.04
8192	23	3.36	4.07	27.32	63.07
16384	25	6.90	9.17		
32768	28	13.83	20.05		
65536	24	27.21	44.85		
131072	27	54.50	104.01		

n	HSS rank A	err TelFun	err CKM	err Rat
512	22	6.67e-14	2.02e-09	1.87e-09
1024	20	1.32e-13	2.70e-09	6.52e-09
2048	21	6.00e-11	3.64e-09	3.73e-09
4096	21	1.99e-13	3.39e-09	4.34e-09
8192	23	1.11e-13	3.72e-09	5.52e-09

factorization, which—as we have seen in section 6.1—is not an optimal choice, partly explaining the observed differences between **expm** and **TelFun**.

6.3. Computation of the inverse square root. To test the presented algorithm for the computation of the inverse square root of an HSS matrix, we consider the problem of sampling from a Gaussian Markov random field (see [6, section 4.2]), which reduces to the computation of the inverse of the square root of a banded matrix. In Table 3 we compare our algorithm with optimal poles, with **CKM** with extended poles (i.e., alternating 0 and ∞); the latter choice of poles is the one made by the authors of **CKM** for solving the presented problem: since the algorithm needs to solve possibly large linear systems, the choice of using mutually different poles can often not be the most advantageous strategy. The number of poles to employ in our method is given by (5.3) (which is in practice very pessimistic), and the accuracy is only used in the determination of the deflation tolerance. The termination of **CKM** is due to the built-in stopping criteria. For completeness, we also approximate $f(A)$ by explicitly evaluating a rational approximation of f : the poles and the residuals of the rational approximation have been derived using the AAA algorithm [25], and for the evaluation, the HSS structure has been exploited using the **hm-toolbox** [23]. In all the cases reported, the degree of the rational approximant constructed by AAA is 12. While **TelFun** is still faster than **CKM** for sufficiently large n , thanks to reduced complexity, its advantage in terms of speed is less evident for this example. Note, however, that its error is significantly lower.

We also show a comparison between **TelFun** and **CKM**, both with (quasi-)optimal poles, for the approximation of the fractional Laplacian, i.e., the computation of $(-A)^{-1/2}$, where A is defined in (6.1). In Figure 4 we compare the timing and the relative error between the presented algorithm and **CKM** using in both cases 50 optimal poles and varying the size of A .




FIG. 4. Comparison of the newly proposed algorithm **TelFun** with **CKM**, for the computation of $f(A)$ where $f(x) = 1/\sqrt{x}$, and A is the discretization of the Laplacian. In both algorithms 50 quasi-optimal poles have been employed.

6.4. Computation of the sign function. In this section, we compute $f(A)$, where f is the sign function, i.e.,

$$f(x) = \begin{cases} 1, & x > 0, \\ -1, & x \leq 0. \end{cases}$$

Assuming that A has both positive and negative eigenvalues (otherwise the computation of $f(A)$ is trivial) the discontinuity of the function does not allow for a reasonable rational approximation on an interval containing the eigenvalues of A . In particular, our convergence result from Theorem 4.5 does not apply. On the other hand, if the eigenvalues of A are contained in $\mathbb{E} = [-b, -a] \cup [a, b]$, with $a, b, > 0$, then the best rational approximation of the sign function on $\mathbb{E}$ is explicitly known in terms of elliptic functions; see [27, section 4.3]. In Table 4, we test the time and the accuracy of the proposed method on tridiagonal matrices whose positive eigenvalues are logarithmically distributed in the interval $[10^a, 1]$ and the negative ones are given by the symmetrization with respect to the imaginary axis. We compare the results with the ones obtained by running **CKM** with optimal poles and with the evaluation of the rational approximation given by the AAA algorithm [25], using the routines contained in the **hm-toolbox** [23].

While not covered by the theory, **TelFun** is clearly the best method and attains good accuracy until $a = -7$. For $a = -9$, the accuracy of all methods suffers from the fact that the eigenvalues get too close to zero.

TABLE 4

Computation of $\text{sign}(A)$, where A is a tridiagonal matrix of size 4096 with logarithmically spaced eigenvalues, symmetric with respect to the imaginary axis contained in $[-1, -10^a] \cup [10^a, 1]$.

a	time TelFun	time CKM	time Rat	err TelFun	err CKM	err Rat
-1	1.73	20.55	50.28	3.75e-10	3.72e-10	1.73e-09
-3	4.13	38.46	128.93	2.60e-10	4.10e-08	1.51e-08
-5	9.72	57.99	121.76	2.70e-10	2.12e-06	7.49e-08
-7	18.24	78.37	137.72	2.99e-08	1.79e-08	1.51e-05
-9	14.08	43.04	139.65	7.45e-02	7.83e-02	3.98e-02




FIG. 5. Comparison of the error behavior in approximating the square root of the discretized Laplacian using the presented algorithm with the bound (4.8), as the number k of poles increases.

6.5. Sharpness of Theorem 4.5. In this section we assess the sharpness of the bound (4.8); to this aim, we consider the matrix $A \in \mathbb{R}^{1024 \times 1024}$ defined in (6.1), and the computation of the square root. We then determine the best rational approximation of the square root over the interval $[\lambda_{\min}, \lambda_{\max}]$, where $\lambda_{\min}$ and $\lambda_{\max}$ are the minimum and maximum eigenvalues of A , respectively, using the `rkfun` gallery [4] for MATLAB. Figure 5 illustrates the error behavior when applying the algorithm `TelFun` to compute the square root of A , utilizing the poles from the best rational approximation. This is compared with the bound provided by (4.8) as the number of poles increases. From this comparison, we can infer that Theorem 4.5 accurately describes the error behavior of the presented algorithm.

7. Conclusions. Generalizing the definition of telescopic decompositions, we have linked different representations of HSS matrices used in the literature, providing ways to convert between them. Exploiting the nested low-rank structure of telescopic decompositions, we have developed a novel algorithm that computes an approximation of $f(A)$ for a symmetric HSS matrix A . Our convergence results imply nearly linear complexity for matrix exponentials and linear-polylogarithmic complexity for inverse square roots in situations of practical relevance. This favorable complexity is attained by using rational Krylov subspaces that involve small-sized matrices only, avoiding the solution of potentially large linear systems usually associated with rational Krylov subspace techniques. Several numerical experiments show that our newly proposed algorithm is faster than existing algorithms for a variety of examples previously reported in the literature. Somewhat surprisingly, it even appears to be the method of choice for computing matrix inverses. A number of questions remain open.

This includes the extension to nonsymmetric matrices as well as a theoretical explanation of the good results obtained for the sign function.

Appendix A. Proofs of Propositions 3.2 and 3.4.

Proof of Proposition 3.2. We proceed by induction on L . If $L = 0$, the tree only consists of the root γ , and the statement trivially holds because of $D_\gamma = A$.

Let us now assume that $L \geq 1$ and consider the matrices $U_\tau^{(\text{big})}$ and $V_\tau^{(\text{big})}$ from Definition 2.1. Because of the recursion (2.1), we have

$$\mathbf{U}^{(L)}(\mathbf{U}^{(L)})^T \begin{bmatrix} U_\alpha^{(\text{big})} & \\ & U_\beta^{(\text{big})} \end{bmatrix} = \begin{bmatrix} U_\alpha^{(\text{big})} & \\ & U_\beta^{(\text{big})} \end{bmatrix},$$

and

$$\mathbf{V}^{(L)}(\mathbf{V}^{(L)})^T \begin{bmatrix} V_\alpha^{(\text{big})} & \\ & V_\beta^{(\text{big})} \end{bmatrix} = \begin{bmatrix} V_\alpha^{(\text{big})} & \\ & V_\beta^{(\text{big})} \end{bmatrix},$$

where α, β are the children of the root γ and $\mathbf{U}^{(L)}, \mathbf{V}^{(L)}$ are the block diagonal matrices employed in Definition 3.1. Combined with the HSS recursion (2.2), this shows that

$$(A.1) \quad A - \begin{bmatrix} A_{\alpha,\alpha} & \\ & A_{\beta,\beta} \end{bmatrix} = \mathbf{U}^{(L)}(\mathbf{U}^{(L)})^T \left(A - \begin{bmatrix} A_{\alpha,\alpha} & \\ & A_{\beta,\beta} \end{bmatrix} \right) \mathbf{V}^{(L)}(\mathbf{V}^{(L)})^T.$$

Noting that $A_{\alpha,\alpha}$ and $A_{\beta,\beta}$ are HSS matrices associated with trees of depth $L-1$, we can apply induction to conclude that they are both in telescopic decomposition. This means that relations of the form (3.1)–(3.2) hold for both matrices or, equivalently,

$$\begin{bmatrix} A_{\alpha,\alpha} & \\ & A_{\beta,\beta} \end{bmatrix} - \mathbf{D}^{(L)} = \mathbf{U}^{(L)}(\mathbf{U}^{(L)})^T \left(\begin{bmatrix} A_{\alpha,\alpha} & \\ & A_{\beta,\beta} \end{bmatrix} - \mathbf{D}^{(L)} \right) \mathbf{V}^{(L)}(\mathbf{V}^{(L)})^T.$$

Adding this equation to (A.1) gives

$$A = \mathbf{D}^{(L)} + \mathbf{U}^{(L)} A^{(L-1)} (\mathbf{V}^{(L)})^T, \quad \text{where} \quad A^{(L-1)} = (\mathbf{U}^{(L)})^T (A - \mathbf{D}^{(L)}) \mathbf{V}^{(L)}.$$

Together with the discussion from section 2.1, it follows that $A^{(L-1)}$ is an HSS matrix associated with $\mathcal{T}_{2r}^{(L-1)}$ and defined by the data-sparse representation

$$\{U_\tau, V_\tau : \tau \in \mathcal{T}_{2r}^{(L-1)}\}, \quad \{\tilde{A}_{\tau,\tau'} : \tau, \tau' \text{ siblings in } \mathcal{T}_{2r}^{(L-1)}\}, \quad \{A_{\tau,\tau}^{(L-1)} : \tau \text{ leaf in } \mathcal{T}_{2r}^{(L-1)}\}.$$

If $A_{\tau,\tau}^{(L-1)} = D_\tau$, this completes the proof by induction: the assumptions of this theorem are satisfied for the level- $(L-1)$ HSS matrix $A^{(L-1)}$ and thus $\{U_\tau, V_\tau, D_\tau\}_{\text{depth}(\tau) \leq L-1}$ is a telescopic decomposition of $A^{(L-1)}$. In turn, all conditions of Definition 3.1 are satisfied and $\{U_\tau, V_\tau, D_\tau\}$ is a telescopic decomposition for A .

It remains to show that $A_{\tau,\tau}^{(L-1)} = D_\tau$ holds for any node τ of depth $L-1$. For this purpose, let α, β denote its children, which are leaves in $\mathcal{T}$. Using that $D_\alpha = A_{\alpha,\alpha}$ and $D_\beta = A_{\beta,\beta}$, we indeed obtain that

$$A_{\tau,\tau}^{(L-1)} = \begin{bmatrix} U_\alpha^T & \\ & U_\beta^T \end{bmatrix} \left(A_{\tau,\tau} - \begin{bmatrix} D_\alpha & \\ & D_\beta \end{bmatrix} \right) \begin{bmatrix} V_\alpha & \\ & V_\beta \end{bmatrix} = \begin{bmatrix} 0 & U_\alpha^T A_{\alpha,\beta} V_\beta \\ U_\beta^T A_{\beta,\alpha} V_\alpha & 0 \end{bmatrix} = D_\tau,$$

where the last equality follows from Definition 2.1. $\square$

Proof of Proposition 3.4. We proceed by induction on L . For $L = 0$, the result trivially holds. Suppose now that $L \geq 1$. According to Definition 3.1, the matrix generated by the standard telescopic decomposition $\{U_\tau, V_\tau, D_\tau\}_{\tau \in \mathcal{T}_{2r}^{(L-1)}}$ is the matrix $A^{(L-1)}$ defined in (3.2). Therefore, if τ is a node of depth $L - 1$ with children α and β , it follows from (3.3) that

$$D_\tau = A_{\tau,\tau}^{(L-1)} = \begin{bmatrix} 0 & U_\alpha^T A_{\alpha,\beta} V_\beta \\ U_\beta^T A_{\beta,\alpha} V_\alpha & 0 \end{bmatrix} = \begin{bmatrix} 0 & \tilde{A}_{\alpha,\beta} \\ \tilde{A}_{\beta,\alpha} & 0 \end{bmatrix},$$

where we set $\tilde{A}_{\alpha,\beta} = U_\alpha^T A_{\alpha,\beta} V_\beta$ and $\tilde{A}_{\beta,\alpha} = U_\beta^T A_{\beta,\alpha} V_\alpha$. This proves (3.4).

It remains to establish the HSS property of A , that is, point 1 of Definition 2.1 (note that point 2 is satisfied by construction). Combining (3.4) with the telescopic relation (3.1), we obtain that

$$A_{\tau,\tau} = \begin{bmatrix} D_\alpha & \\ & D_\beta \end{bmatrix} + \begin{bmatrix} U_\alpha & \\ & U_\beta \end{bmatrix} D_\tau \begin{bmatrix} V_\alpha & \\ & V_\beta \end{bmatrix} = \begin{bmatrix} D_\alpha & U_\alpha \tilde{A}_{\alpha,\beta} V_\beta^T \\ U_\beta \tilde{A}_{\beta,\alpha} V_\alpha^T & D_\beta \end{bmatrix}.$$

In particular, $A_{\alpha,\beta} = U_\alpha \tilde{A}_{\alpha,\beta} V_\beta^T$, which establishes point 1 of Definition 2.1 for two sibling leaves $\tau = \alpha$, $\tau' = \beta$. To show the corresponding property for two siblings τ and τ' of depth $\ell < L$, let $\alpha_1, \dots, \alpha_{2^{L-\ell}}$ and $\alpha'_1, \dots, \alpha'_{2^{L-\ell}}$ denote the leaf nodes in the corresponding subtrees. Using (3.1), the proof is completed by noting that

$$\begin{aligned} A_{\tau,\tau'} &= \text{blkdiag}(U_{\alpha_i}) A_{\tau,\tau'}^{(L-1)} \text{blkdiag}(V_{\alpha'_i})^T \\ &= \text{blkdiag}(U_{\alpha_i}) \text{blkdiag}(U_{\alpha_i})^T U_\tau^{(\text{big})} \tilde{A}_{\tau,\tau'} (V_{\tau'}^{(\text{big})})^T \text{blkdiag}(V_{\alpha'_i}) \text{blkdiag}(V_{\alpha'_i})^T \\ &= U_\tau^{(\text{big})} \tilde{A}_{\tau,\tau'} (V_{\tau'}^{(\text{big})})^T, \end{aligned}$$

where the second equality uses induction: $A^{(L-1)}$ is an HSS matrix and satisfies a relation of the form (2.2) for the parent γ of τ, τ' . $\square$

Acknowledgments. The authors thank the reviewers for their careful reading and constructive feedback. Parts of this work were prepared while the first author was visiting the ANCHP group at EPFL. Furthermore, during part of the preparation of this work, the first author was affiliated with the Scuola Normale Superiore in Pisa, Italy. The first and third authors are both members of the research group INdAM-GNCS.

REFERENCES

- [1] B. BECKERMANN, A. CORTINOVIS, D. KRESSNER, AND M. SCHWEITZER, *Low-rank updates of matrix functions II: Rational Krylov methods*, SIAM J. Numer. Anal., 59 (2021), pp. 1325–1347, <https://doi.org/10.1137/20M1362553>.
- [2] B. BECKERMANN AND L. REICHEL, *Error estimates and evaluation of matrix functions via the Faber transform*, SIAM J. Numer. Anal., 47 (2009), pp. 3849–3883, <https://doi.org/10.1137/080741744>.
- [3] M. BENZI AND V. SIMONCINI, *Decay bounds for functions of Hermitian matrices with banded or Kronecker structure*, SIAM J. Matrix Anal. Appl., 36 (2015), pp. 1263–1282, <https://doi.org/10.1137/151006159>.
- [4] M. BERLJAFI AND S. GÜTTTEL, *The RKFIT algorithm for nonlinear rational approximation*, SIAM J. Sci. Comput., 39 (2017), pp. A2049–A2071, <https://doi.org/10.1137/15M1025426>.
- [5] S. CHANDRASEKARAN, M. GU, AND T. PALS, *A fast ULV decomposition solver for hierarchically semiseparable representations*, SIAM J. Matrix Anal. Appl., 28 (2006), pp. 603–622, <https://doi.org/10.1137/S0895479803436652>.

- [6] A. CORTINOVIS, D. KRESSNER, AND S. MASSEI, *Divide-and-conquer methods for functions of matrices with banded or hierarchical low-rank structure*, SIAM J. Matrix Anal. Appl., 43 (2022), pp. 151–177, <https://doi.org/10.1137/21M1432594>.
- [7] S. ELSWORTH AND S. GÜTTEL, *The block rational Arnoldi method*, SIAM J. Matrix Anal. Appl., 41 (2020), pp. 365–388, <https://doi.org/10.1137/19M1245505>.
- [8] A. FROMMER, C. SCHIMMEL, AND M. SCHWEITZER, *Analysis of probing techniques for sparse approximation and trace estimation of decaying matrix functions*, SIAM J. Matrix Anal. Appl., 42 (2021), pp. 1290–1318, <https://doi.org/10.1137/20M1364461>.
- [9] A. GILLMAN, P. M. YOUNG, AND P.-G. MARTINSSON, *A direct solver with $O(N)$ complexity for integral equations on one-dimensional domains*, Front. Math. China, 7 (2012), pp. 217–247, <https://doi.org/10.1007/s11464-012-0188-3>.
- [10] A. A. GONCHAR AND E. A. RAKHMANOV, *Equilibrium distributions and the rate of rational approximation of analytic functions*, Mat. Sb. (NS), 134 (1987), pp. 306–352, 447.
- [11] L. GRASEDYCK, W. HACKBUSCH, AND B. N. KHOROMSKIY, *Solution of large scale algebraic matrix Riccati equations by use of hierarchical matrices*, Computing, 70 (2003), pp. 121–165, <https://doi.org/10.1007/s00607-002-1470-0>.
- [12] S. GÜTTEL, *Rational Krylov Methods for Operator Functions*, Ph.D. thesis, Technische Universität Bergakademie Freiberg, 2010.
- [13] S. GÜTTEL, *Rational Krylov approximation of matrix functions: Numerical methods and optimal pole selection*, GAMM-Mitt., 36 (2013), pp. 8–31, <https://doi.org/10.1002/gamm.201310002>.
- [14] W. HACKBUSCH, *Hierarchical Matrices: Algorithms and Analysis*, Springer Series in Computational Mathematics 49, Springer, Heidelberg, 2015, <https://doi.org/10.1007/978-3-662-47324-5>.
- [15] D. HALIKIAS AND A. TOWNSEND, *Structured matrix recovery from matrix-vector products*, Numer. Linear Algebra Appl., 31 (2024), e2531, <https://doi.org/10.1002/nla.2531>.
- [16] N. J. HIGHAM, *Functions of Matrices*, SIAM, Philadelphia, 2008, <https://doi.org/10.1137/1.9780898717778>.
- [17] N. J. HIGHAM, *The scaling and squaring method for the matrix exponential revisited*, SIAM Rev., 51 (2009), pp. 747–764, <https://doi.org/10.1137/090768539>.
- [18] M. HOCHBRUCK AND C. LUBICH, *On Krylov subspace approximations to the matrix exponential operator*, SIAM J. Numer. Anal., 34 (1997), pp. 1911–1925, <https://doi.org/10.1137/S0036142995280572>.
- [19] D. KRESSNER AND A. ŠUŠNJARA, *Fast computation of spectral projectors of banded matrices*, SIAM J. Matrix Anal. Appl., 38 (2017), pp. 984–1009, <https://doi.org/10.1137/16M1087278>.
- [20] J. LEVITT AND P.-G. MARTINSSON, *Linear-complexity black-box randomized compression of rank-structured matrices*, SIAM J. Sci. Comput. 46 (2024), <https://doi.org/10.1137/22M1528574>.
- [21] P.-G. MARTINSSON, *A fast randomized algorithm for computing a hierarchically semiseparable representation of a matrix*, SIAM J. Matrix Anal. Appl., 32 (2011), pp. 1251–1274, <https://doi.org/10.1137/100786617>.
- [22] S. MASSEI, M. MAZZA, AND L. ROBOL, *Fast solvers for two-dimensional fractional diffusion equations using rank structured matrices*, SIAM J. Sci. Comput., 41 (2019), pp. A2627–A2656, <https://doi.org/10.1137/18M1180803>.
- [23] S. MASSEI, L. ROBOL, AND D. KRESSNER, *hm-toolbox: MATLAB software for HODLR and HSS matrices*, SIAM J. Sci. Comput., 42 (2020), pp. C43–C68, <https://doi.org/10.1137/19M1288048>.
- [24] M. M. MEERSCHAERT AND C. TADJERAN, *Finite difference approximations for fractional advection-dispersion flow equations*, J. Comput. Appl. Math., 172 (2004), pp. 65–77, <https://doi.org/10.1016/j.cam.2004.01.033>.
- [25] Y. NAKATSUKASA, O. SÈTE, AND L. N. TREFETHEN, *The AAA algorithm for rational approximation*, SIAM J. Sci. Comput., 40 (2018), pp. A1494–A1522, <https://doi.org/10.1137/16M1106122>.
- [26] T. PARK AND Y. NAKATSUKASA, *Approximating Sparse Matrices and Their Functions Using Matrix-Vector Products*, preprint, arXiv:2310.05625, 2023.
- [27] P. P. PETRUSHEV AND V. A. POPOV, *Rational Approximation of Real Functions*, Encyclopedia of Mathematics and Its Applications 28, Cambridge University Press, Cambridge, 1987.
- [28] J. XIA, S. CHANDRASEKARAN, M. GU, AND X. S. LI, *Fast algorithms for hierarchically semiseparable matrices*, Numer. Linear Algebra Appl., 17 (2010), pp. 953–976, <https://doi.org/10.1002/nla.691>.