

MATHICSE Technical Report

Nr. 31 .2016

August 2016



Fast computation of spectral
projectors of banded matrices

Daniel Kressner, Ana Šušnjara

Fast computation of spectral projectors of banded matrices

Daniel Kressner* Ana Šušnjara†

Abstract

We consider the approximate computation of spectral projectors for symmetric banded matrices. While this problem has received considerable attention, especially in the context of linear scaling electronic structure methods, the presence of small relative spectral gaps challenges existing methods based on approximate sparsity. In this work, we show how a data-sparse approximation based on hierarchical matrices can be used to overcome this problem. We prove a priori bounds on the approximation error and propose a fast algorithm based on the QDWH algorithm, along the works by Nakatsukasa et al. Numerical experiments demonstrate that the performance of our algorithm is robust with respect to the spectral gap. A preliminary MATLAB implementation becomes faster than `eig` already for matrix sizes of a few thousand.

1 Introduction

Given a symmetric banded matrix $A \in \mathbb{R}^{n \times n}$ with eigenvalues

$$\lambda_1 \leq \dots \leq \lambda_\nu < \mu < \lambda_{\nu+1} \leq \dots \leq \lambda_n,$$

we consider the computation of the spectral projector $\Pi_{<\mu}(A)$ associated with the eigenvalues $\lambda_1, \dots, \lambda_\nu$. We specifically target the situation where both n and ν are large, say $n = 100\,000$ and $\nu = 50\,000$, which makes approaches based on computing eigenvectors computationally expensive. For a tridiagonal matrix, the MRRR algorithm requires $\mathcal{O}(\nu n)$ operations and memory [18] to compute the ν eigenvectors needed to define $\Pi_{<\mu}(A)$.

There are a number of applications giving rise to the problem under consideration. First and foremost, this task is at the heart of linear scaling methods for the calculation of the electronic structure of molecules with a large number of atoms. For insulators at zero temperature, the density matrix is the spectral projector associated with the eigenvalues of the Hamiltonian below the so called HOMO-LUMO gap; see [23] for an overview. The Hamiltonian is usually symmetric and, depending on the discretization and the structure of the molecule, it can be (approximately) banded. A number of existing linear scaling methods use that this sometimes implies that the spectral projector may also admit a good approximation by a banded matrix; see [10] for a recent survey and a mathematical justification. For this approach to work well, the HOMO-LUMO gap should not become too small. For metallic systems, this gap actually converges to zero, which

*MATHICSE-ANCHP, École Polytechnique Fédérale de Lausanne, Station 8, 1015 Lausanne, Switzerland. E-mail: daniel.kressner@epfl.ch.

†MATHICSE-ANCHP, École Polytechnique Fédérale de Lausanne, Station 8, 1015 Lausanne, Switzerland. E-mail: ana.susnjara@epfl.ch. The work of Ana Šušnjara has been supported by the SNSF research project *Low-rank updates of matrix functions and fast eigenvalue solvers*.

makes it impossible to apply an approach based on approximate bandedness or, more generally, sparsity.

Another potential important application for banded matrices arises in dense symmetric eigenvalue solvers. The eigenvalues and eigenvectors of a symmetric dense matrix A are usually computed by first reducing A to tridiagonal form and then applying either divide-and-conquer method or MRRR; see, e.g. [4, 17] for recent examples. It is by no means trivial to implement the reduction to tridiagonal form efficiently so that it performs well on a modern computing architecture with a memory hierarchy. Most existing approaches [3, 12, 29, 31, 43], with the notable exception of [41], are based on successive band reduction [13]. In this context, it would be preferable to design an eigenvalue solver that works directly with banded matrices, bypassing the need for tridiagonal reduction. While we are not aware of any such extension of MRRR, this possibility has been explored several times for the divide-and-conquer method, e.g., in [2, 30]. The variants proposed so far seem to suffer either from numerical instabilities or from a complexity that grows significantly with the bandwidth. The method proposed in this paper can be used to directly compute the spectral projector of a banded matrix, which in turn could potentially be used as a basis for a fast spectral divide and conquer algorithm in the spirit of Nakatsukasa and Higham [39].

To deal with small spectral gaps, one needs to go beyond sparsity. It turns out that hierarchical matrices [27], also called \mathcal{H} -matrices, are much better suited in such a setting. Intuitively, this can be well explained by considering the approximation of the Heaviside function $\Pi_{<\mu}(x)$ on the eigenvalues of A . While a polynomial approximation of $\Pi_{<\mu}$ corresponds to a sparse approximation of $\Pi_{<\mu}(A)$ [10], a rational approximation corresponds to an approximation of $\Pi_{<\mu}(A)$ that features hierarchical low-rank structure. It is well known, see, e.g., [40], that a rational approximation is more powerful in dealing with nearby singularities, such as $x = \mu$ for $\Pi_{<\mu}(x)$.

There are a number of existing approaches to use hierarchical low-rank structures for the fast computation of matrix functions, including spectral projectors. Beylkin, Coult, and Mohlenkamp [11] proposed a combination of the Newton–Schulz iteration with the HODLR format, a subset of \mathcal{H} -matrices, to compute spectral projectors for banded matrices. However, the algorithm does not fully exploit the potential of low-rank formats; it converts a full matrix to the HODLR format in each iteration. In the context of Riccati and Lyapunov matrix equations, the computation of the closely related sign function of an \mathcal{H} -matrix has been discussed in [25, 5]. The work in [21, 22, 25] involves the \mathcal{H} -matrix approximation of resolvents, which is then used to compute the matrix exponential and related matrix functions.

Other hierarchical matrix techniques for eigenvalue problems include slicing-the-spectrum, which uses LDL decompositions to compute eigenvalues in a specified interval for symmetric HODLR and HSS matrices [9] as well as \mathcal{H}^2 -matrices [7]. Approximate \mathcal{H} -matrix inverses can be used as preconditioners in iterative eigenvalue solvers; see [33, 35] for examples. Recently, Vogel et al. [45] have developed a fast divide-and-conquer method for computing all eigenvalues and eigenvectors in the HSS format. However, as the matrix of eigenvectors is represented in a factored form, it would be a nontrivial and possibly expensive detour to compute spectral projectors via this approach.

In this paper we propose a new method based on a variant [39] of the QR-based dynamically weighted Halley algorithm (QDWH) for computing a polar decomposition [37]. Our method exploits the fact that the iterates of QDWH applied to a banded matrix can be well approximated in the HODLR format. In fact, we show that the memory needed for storing the approximate spectral projector depends only logarithmically on the spectral gap, a major improvement over approximate sparsity. The implementation of QDWH requires some care, in particular, concerning the representation of the first iterate. One major contribution of this work is to show how

this can be done efficiently.

The remainder of the paper is organized as follows. In Section 2, we review the QDWH algorithm for computing a spectral projector $\Pi_{<\mu}(A)$. Section 3 recalls well-known facts about the HODLR format and the corresponding formatted arithmetics. Based on the best rational approximation to the sign function, we derive new a priori bounds on the singular values for off-diagonal blocks of $\Pi_{<\mu}(A)$, from which we deduce bounds on the memory required to store $\Pi_{<\mu}(A)$ approximately in the HODLR format. Section 4 discusses the efficient realization of the QR decomposition required in the first iterate of the QDWH algorithm. Section 5 summarizes our newly proposed QDWH algorithm in the HODLR format and provides implementation details. Finally, numerical experiments both for tridiagonal and banded matrices are shown in Section 6.

2 Computation of spectral projectors via QDWH

In the following, we assume $\mu = 0$ without loss of generality, and thus consider the computation of the spectral projector $\Pi_{<0}(A)$ associated with the negative eigenvalues of a symmetric non-singular matrix $A \in \mathbb{R}^{n \times n}$. Following [39], our approach is based on a well-known connection to the polar decomposition.

The polar decomposition [24, Chapter 9] of A takes the form $A = UH$ for an orthogonal matrix U and a symmetric positive definite matrix H . Let $A = V\Lambda V^*$ be a spectral decomposition of A such that $\Lambda = \text{diag}(\Lambda_-, \Lambda_+)$, where Λ_- and Λ_+ are diagonal matrices containing the ν negative and the $n - \nu$ positive eigenvalues of A , respectively. Then

$$\begin{aligned} A &= V \text{diag}(\Lambda_-, \Lambda_+) V^* \\ &= \underbrace{V \text{diag}(-I_\nu, I_{n-\nu}) V^*}_{=:U} \cdot \underbrace{V \text{diag}(|\Lambda_-|, |\Lambda_+|) V^*}_{=:H} \end{aligned}$$

gives the polar decomposition of A . In particular, this shows that the matrix sign function $\text{sign}(A)$ coincides with the orthogonal factor U from the polar decomposition. More importantly, $\Pi_{<0}(A) = \frac{1}{2}(I - U)$.

2.1 QDWH algorithm

The QDWH algorithm [37] computes the polar factor U of A as the limit of the sequence X_k defined by

$$\begin{aligned} X_0 &= A/\alpha, \\ X_{k+1} &= X_k(a_k I + b_k X_k^* X_k)(I + c_k X_k^* X_k)^{-1}. \end{aligned} \quad (1)$$

The parameter $\alpha > 0$ is an estimate of $\|A\|_2$. The parameters a_k, b_k, c_k are computed via the relations

$$a_k = h(l_k), \quad b_k = (a_k - 1)^2/4, \quad c_k = a_k + b_k - 1. \quad (2)$$

Representing a lower bound for the smallest singular value of X_k , the parameter l_k is determined by the recurrence

$$l_k = l_{k-1}(a_{k-1} + b_{k-1}l_{k-1}^2)/(1 + c_{k-1}l_{k-1}^2), \quad k \geq 1,$$

where l_0 is a lower bound for $\sigma_{\min}(X_0)$. The function h is given by

$$h(l) = \sqrt{1 + \gamma} + \frac{1}{2} \sqrt{8 - 4\gamma + \frac{8(2 - l^2)}{l^2 \sqrt{1 + \gamma}}}, \quad \gamma = \sqrt[3]{\frac{4(1 - l^2)}{l^4}}.$$

The efficient estimation of α and l_0 , required to start the recurrence, will be discussed in Section 5.

The QDWH algorithm is cubically convergent and it has been shown in [37] that at most $k = 6$ iterations are needed to obtain convergence within tolerance 10^{-16} , i.e. $\|X_6 - U\|_2 < 10^{-16}$ for every matrix A with $\kappa(A) \leq 10^{16}$.

The recurrence (1) has the equivalent form

$$X_0 = A/\alpha, \quad (3a)$$

$$X_{k+1} = \frac{b_k}{c_k} X_k + \frac{1}{\sqrt{c_k}} \left(a_k - \frac{b_k}{c_k} \right) Q_1 Q_2^*, \quad (3b)$$

with the QR decomposition

$$\begin{bmatrix} \sqrt{c_k} X_k \\ I \end{bmatrix} = \begin{bmatrix} Q_1 \\ Q_2 \end{bmatrix} R. \quad (4)$$

Throughout the paper, we refer to (3) as a *QR-based iteration*. On the other hand, as observed in [39], the recurrence (1) can also be rewritten in terms of the *Cholesky-based iteration*

$$Z_k = I + c_k X_k^* X_k, \quad W_k = \text{chol}(Z_k), \quad (5a)$$

$$X_{k+1} = \frac{b_k}{c_k} X_k + \left(a_k - \frac{b_k}{c_k} \right) (X_k W_k^{-1}) W_k^{-*}, \quad (5b)$$

where $\text{chol}(Z_k)$ denotes the Cholesky factor of Z_k .

Following [37], either variant of the QDWH algorithm is terminated when l_k is sufficiently close to 1, that is, $|1 - l_k| \leq \delta$ for some stopping tolerance δ , say $\delta = 10^{-15}$.

We mention that a higher-order variant of QDWH, called Zolo-pd, has recently been proposed by Freund and Nakatsukasa [38]. This method approximates the polar decomposition in at most two iterations but requires more arithmetic per iteration.

2.2 Switching between QR-based and Cholesky-based iterations

Due to its lower operation count, it can be expected that one Cholesky-based iteration (5) is faster than one QR-based iteration (3). However, when Z_k is ill-conditioned, which is signaled by a large value of c_k , the numerical stability of (5) can be jeopardized. To avoid this, it is proposed in [39] to switch from (3) to (5) as soon as $c_k \leq 100$. Since c_k converges monotonically from above to 3, this implies that this hybrid approach will first perform a few QR-based iterations and then switch for good to Cholesky-based iterations. In fact, numerical experiments presented in [39] indicate that at most two QR-based iterations are performed.

For reasons explained in Remark 2 below, we prefer to perform only one QR-based iteration and then switch to Cholesky-based iterations. To explore the impact of this choice on numerical accuracy, we perform a comparison of the QDWH algorithm proposed in [39] with a variant of QDWH that performs only one QR-based iteration. We consider the following error measures:

$$\begin{aligned} e_{\text{id}}^Q &:= \|U^2 - I\|_2, \\ e_{\text{trace}}^Q &:= |\text{trace}(U) - \text{trace}(\text{sign}(A))|, \\ e_{\text{SP}}^Q &:= \left\| \frac{1}{2}(I - U) - \Pi_{<0}(A) \right\|_2, \end{aligned} \quad (6)$$

where U denotes the output of the QDWH algorithm, and $\Pi_{<0}(A)$ the spectral projector returned by the MATLAB function `eig`.

Algorithm [39]	gap	10^{-1}	10^{-5}	10^{-10}	10^{-15}
one QR-based iteration (3)	e_{trace}^Q	$5.55 \cdot 10^{-17}$	$7.22 \cdot 10^{-16}$	$2.22 \cdot 10^{-16}$	$1.11 \cdot 10^{-16}$
	e_{id}^Q	$1.15 \cdot 10^{-15}$	$2.41 \cdot 10^{-15}$	$1.84 \cdot 10^{-15}$	$1.82 \cdot 10^{-15}$
	e_{SP}^Q	$1.87 \cdot 10^{-14}$	$4.35 \cdot 10^{-12}$	$1.88 \cdot 10^{-6}$	$1.91 \cdot 10^{-2}$
several QR-based iterations (3)	e_{trace}^Q	$5.55 \cdot 10^{-17}$	$1.22 \cdot 10^{-15}$	$1.53 \cdot 10^{-16}$	$6.25 \cdot 10^{-16}$
	e_{id}^Q	$1.15 \cdot 10^{-15}$	$2.58 \cdot 10^{-15}$	$1.81 \cdot 10^{-15}$	$2.04 \cdot 10^{-15}$
	e_{SP}^Q	$1.87 \cdot 10^{-14}$	$2.12 \cdot 10^{-12}$	$2.82 \cdot 10^{-6}$	$3.06 \cdot 10^{-2}$
	# of (3)	1	2	2	3

Table 1: Comparison of errors in the QDWH algorithm with one or several QR-based iterations.

Example 1. Let $A \in \mathbb{R}^{2000 \times 2000}$ be a symmetric tridiagonal matrix constructed as described in Section 6.1, such that half of the spectrum of A is contained in $[-1, -\text{gap}]$ and the other half in $[\text{gap}, 1]$, for $\text{gap} \in \{10^{-1}, 10^{-5}, 10^{-10}, 10^{-15}\}$. As can be seen in Table 1, the errors obtained by both variants of the QDWH algorithm exhibit a similar behavior. Even for tiny spectral gaps, no significant loss of accuracy is observed if only one QR-based iteration is performed.

3 Hierarchical matrix approximation of spectral projectors

Introduced in the context of integral and partial differential equations, hierarchical matrices allow for the data-sparse representation of a certain class of dense matrices. In the following, we briefly recall the concept of hierarchical matrices and some operations; see, e.g., [6, 28] for more details.

3.1 Matrices with hierarchical low-rank structures

3.1.1 HODLR matrices

We first discuss *hierarchically off-diagonal low-rank* (HODLR) matrices. For convenience, we assume that $n = 2^p$ for $p \in \mathbb{N}$. Given a prescribed maximal off-diagonal rank $k \in \mathbb{N}$, we suppose that a matrix $M \in \mathbb{R}^{n \times n}$ admits the representation

$$M = \begin{bmatrix} M_1^{(1)} & U_1^{(1)} V_1^{(1)*} \\ U_2^{(1)} V_2^{(1)*} & M_2^{(1)} \end{bmatrix}, \quad (7)$$

where $M_i^{(1)} \in \mathbb{R}^{\frac{n}{2} \times \frac{n}{2}}$, $U_i^{(1)}, V_i^{(1)} \in \mathbb{R}^{\frac{n}{2} \times k}$, for $i = 1, 2$, and $k \ll n$. A HODLR matrix is obtained by applying (7) recursively to the diagonal blocks $M_i^{(l-1)}$, where $i = 1, \dots, 2^{l-1}$ for the l th level of recursion, $2 \leq l \leq p$. The recursion terminates when the diagonal blocks are sufficiently small, that is, $\frac{n}{2^l} \leq n_{\min}$ for a minimal block size $n_{\min} \in \mathbb{N}$; see Figure 1 below for an illustration. Formally, we define the set of HODLR matrices with block-wise rank k as

$$\mathcal{H}(k) := \{M \in \mathbb{R}^{n \times n} : \text{rank } M|_{\text{off}} \leq k \ \forall \text{off-diagonal block } M|_{\text{off}} \text{ in recursive subdivision}\}.$$

Any matrix $M \in \mathcal{H}(k)$ admits a data-sparse representation. By storing the off-diagonal blocks in terms of their low-rank factors and the diagonal blocks as dense matrices, the memory required for representing M is $\mathcal{O}(kn \log n)$, assuming that k is constant with respect to n .

Given a general matrix $A \in \mathbb{R}^{n \times n}$, an approximation $M \in \mathcal{H}(k)$ to A is obtained by computing truncated singular value decompositions of the off-diagonal blocks of A . The quality of such an

approximation is governed by the truncated singular values. For simplifying the presentation, we have assumed that the ranks in the off-diagonal blocks are all bounded by the same integer k . In practice, we choose these ranks adaptively based on an absolute truncation tolerance ϵ and they may be different for each block.

As explained in [6, 28], several matrix operations can be performed approximately and efficiently within the HODLR format. The use of formatted arithmetics leads to linear-polylogarithmic complexity for these operations. Table 2 summarizes the complexity of operations needed by the QDWH algorithm for $M_1, M_2, R \in \mathcal{H}(k)$, where T is triangular, and $v \in \mathbb{R}^n$.

Table 2: Complexity of some arithmetic operations in the HODLR format.

Operation	Computational complexity
Matrix-vector mult. $M_1 *_{\mathcal{H}} v$	$\mathcal{O}(kn \log n)$
Matrix addition $M_1 +_{\mathcal{H}} M_2 \in \mathcal{H}(k)$	$\mathcal{O}(k^2 n \log n)$
Matrix multiplication $M_1 *_{\mathcal{H}} M_2 \in \mathcal{H}(k)$	$\mathcal{O}(k^2 n \log^2 n)$
Cholesky decomposition \mathcal{H} -Cholesky(M_1) $\in \mathcal{H}(k)$	$\mathcal{O}(k^2 n \log^2 n)$
Solving triangular system $M_1 *_{\mathcal{H}} T = M_2 \in \mathcal{H}(k)$	$\mathcal{O}(k^2 n \log^2 n)$

Remark 2. *The QR-based iteration (3) of QDWH requires the computation of the QR decomposition (4). Unlike for \mathcal{H} -Cholesky, there is no straightforward way of performing QR decompositions in hierarchical matrix arithmetics. To our knowledge, three different algorithms [6, 8, 34] have been proposed for this purpose. However, each of them seems to have some drawbacks, e.g., failing to achieve a highly accurate decomposition or leading to loss of orthogonality in the orthogonal factor. Hence, instead of using any of the existing algorithms, we develop a novel method in Section 4 to compute the QR decomposition (3) that exploits the particular structure of the matrix in the first iteration of the QDWH algorithm.*

3.1.2 Hierarchical matrices

Let $I = \{1, 2, \dots, n\}$ denote the row and column index sets of a matrix $M \in \mathbb{R}^{n \times n}$. To consider more general hierarchical matrices, we define a partition P of $I \times I$ as follows. On level $l = 0$, the index set $I^0 := I$ is partitioned into $I^0 = I_1^1 \cup I_2^1$, with $I_1^1 = \{1, \dots, \frac{n}{2}\}$ and $I_2^1 = \{\frac{n}{2} + 1, \dots, n\}$. At this point, the partition P contains five blocks: $I \times I$ and $I_i^1 \times I_j^1$ for $i, j = 1, 2$. The subdivision continues as follows: on each level $l = 1, \dots, p-1$ the index sets I_i^l are partitioned into sets I_{2i-1}^{l+1} and I_{2i}^{l+1} of equal size, contributing the blocks $I_i^{l+1} \times I_j^{l+1}$ for $i, j = 1, \dots, 2^l$ to the partition P . The recursion terminates when a block $I_i^l \times I_j^l$ satisfies a certain admissibility condition or when $\min\{|I_i^l|, |I_j^l|\} \leq n_{\min}$ holds.

Inspired by discretizations for 1D integral equations [27], we make use of the following admissibility condition:

$$\text{block } \tau = t \times s \text{ is admissible} \iff \min\{\text{diam}(t), \text{diam}(s)\} \leq \text{dist}(t, s), \quad (8)$$

with

$$\text{diam}(t) := \max_{i, j \in t} |i - j|, \quad \text{dist}(t, s) := \min_{i \in t, j \in s} |i - j|.$$

See Figure 1 for an illustration of the resulting partition P . Given P , the set of \mathcal{H} -matrices with block-wise rank k is defined as

$$\mathcal{H}(P, k) := \{M \in \mathbb{R}^{n \times n} : \text{rank } M|_{\tau} \leq k \text{ for all admissible blocks } \tau \in P\}.$$

The complexity of arithmetic operations displayed in Table 2 extends to $\mathcal{H}(P, k)$.




Figure 1: Left: HODLR matrix. Right: \mathcal{H} -matrix with admissibility condition (8). Blocks colored grey are stored as dense matrices.

Example 3. We investigate the potential of the HODLR and \mathcal{H} -matrix formats to efficiently store spectral projectors of banded matrices. For this purpose, we have generated, as explained in Section 6.1, a symmetric b -banded matrix $A \in \mathbb{R}^{16000 \times 16000}$ with eigenvalues in $[-1, -\text{gap}] \cup [\text{gap}, 1]$. The memory needed to store the full spectral projector $\Pi_{<0}(A)$ in double precision is 2048 MB. We choose $n_{\min} = 250$, a truncation tolerance $\epsilon = 10^{-10}$, and $\text{gap} \in \{10^{-1}, 10^{-4}\}$. Table 3 reveals that the HODLR format often requires less memory to approximately store $\Pi_{<0}(A)$, unless both gap and the bandwidth are large. In terms of computational time, the outcome is even clearer. For bandwidth $b = 8$ and $\text{gap} = 10^{-1}$, a situation that favors the \mathcal{H} -matrix format in terms of memory, we have run the algorithm described in Section 5 in both formats. It turned out that the use of the HODLR format led to an overall time of 608 seconds, while the \mathcal{H} -matrix format required 792 seconds.

Table 3: Memory required to approximately store spectral projectors for the banded matrices from Example 3 in HODLR and \mathcal{H} -matrix format.

$\text{gap} = 10^{-1}$	HODLR	\mathcal{H} -matrix	$\text{gap} = 10^{-4}$	HODLR	\mathcal{H} -matrix
b = 1	55.72 MB	95.16 MB	b = 1	86.03 MB	128.58 MB
b = 2	79.38 MB	96.42 MB	b = 2	129.71 MB	160.56 MB
b = 4	127.04 MB	106.54 MB	b = 4	206.32 MB	225.72 MB
b = 8	219.92 MB	151.06 MB	b = 8	340.88 MB	352.54 MB
b = 16	395.91 MB	291.85 MB	b = 16	567.69 MB	583.93 MB

Based on the evidence provided by Example 3, we have concluded that more general \mathcal{H} -matrix formats bring little advantage and thus focus on the HODLR format for the rest of this paper.

3.2 A priori bounds on singular values and memory requirements

To study the approximation of $\Pi_{<0}(A)$ in the HODLR format, we first derive bounds for the singular values of the off-diagonal ranks based on rational approximations to the sign function. In the following, we say that a rational function r is of type (k, s) and write $r \in \mathcal{R}_{k,s}$ if $r = p/q$ holds for polynomials p and q of degree at most k and s , respectively.

3.2.1 Rational approximation of sign function

Given $R > 0$, the min-max problem

$$\min_{r \in \mathcal{R}_{2m-1, 2m}} \max_{x \in [-R, -1] \cup [1, R]} |\text{sign}(x) - r(x)| \quad (9)$$

has a unique solution s_m . Called a Zolotarev function of type $(2m-1, 2m)$ corresponding to R (see e.g. [1, Chapter 9]), this function takes the form

$$s_m(x) := Cx \frac{\prod_{i=1}^{m-1} (x^2 + c_{2i})}{\prod_{i=1}^m (x^2 + c_{2i-1})}.$$

The coefficients $c_i, i = 1, \dots, 2m$ are given in terms of the Jacobi elliptic function $\text{sn}(\cdot; \kappa)$:

$$c_i = \frac{\text{sn}^2\left(\frac{iK(\kappa)}{2m}; \kappa\right)}{1 - \text{sn}^2\left(\frac{iK(\kappa)}{2m}; \kappa\right)}, \quad (10)$$

where $\kappa = \sqrt{1 - 1/R^2}$ and $K(\kappa)$ is defined as the complete elliptic integral of the first kind

$$K(\kappa) = \int_0^{\frac{\pi}{2}} \frac{d\theta}{\sqrt{1 - \kappa^2 \sin^2 \theta}} = \int_0^1 \frac{dt}{\sqrt{(1-t^2)(1-\kappa^2 t^2)}}.$$

The constant C is uniquely determined by the condition

$$\min_{x \in [-R, -1]} 1 + s_m(x) = \max_{x \in [1, R]} 1 - s_m(x).$$

As shown in [26], the approximation error $E_m := \max_{x \in [-R, -1] \cup [1, R]} |\text{sign}(x) - s_m(x)|$ is bounded as

$$\frac{4\rho^m}{\rho^m + 1} \leq E_m \leq 4\rho^m, \quad (11)$$

where $\rho = \rho(\mu) = \exp\left(-\frac{\pi K(\mu')}{2K(\mu)}\right)$ with $\mu = \left(\frac{\sqrt{R}-1}{\sqrt{R}+1}\right)^2$ and $\mu' = \sqrt{1 - \mu^2}$. The following lemma derives a bound from (11) that reveals the influence of the gap on the error.

Lemma 4. *With the notation introduced above and $\text{gap} = 1/R$, it holds that*

$$E_m \leq 4 \exp\left(-\frac{\pi^2 m}{4 \log(4/\sqrt[4]{\text{gap}} + 2)}\right). \quad (12)$$

Proof. Following Braess and Hackbusch [15], we have

$$K(\mu') \geq \pi/2, \quad K(\mu) \leq \log(4/\mu' + 2).$$

Thus, the upper bound in (11) implies

$$E_m \leq 4 \exp\left(-\frac{\pi^2 m}{4 \log(4/\mu' + 2)}\right).$$

From

$$\mu' = \sqrt{1 - \left(\frac{1 - \sqrt{\text{gap}}}{1 + \sqrt{\text{gap}}}\right)^4} = \frac{\sqrt{8\sqrt{\text{gap}}(1 + \text{gap})}}{(1 + \sqrt{\text{gap}})^2} \geq \sqrt[4]{\text{gap}}$$

it follows that $\log(4/\mu' + 2) \leq \log(4/\sqrt[4]{\text{gap}} + 2)$, which completes the proof. \square

It is simple to bound the ranks of the off-diagonal blocks for a rational function applied to a banded matrix.

Lemma 5. *Consider a b -banded matrix $A \in \mathbb{R}^{n \times n}$ and a rational function r_m of type $(2m - 1, 2m)$, with poles disjoint from the spectrum of A . Then the off-diagonal blocks of $r_m(A)$ have rank at most $2mb$.*

Proof. Assuming that r has simple poles, let $r_m(x) = \sum_{i=1}^{2m} \omega_i(x - \mu_i)^{-1}$ be a partial fraction expansion of r_m , with $\omega_i, \mu_i \in \mathbb{C}, i = 1, \dots, 2m$. Thus, $r_m(A)$ is a sum of $2m$ shifted inverses of A . By a well known result (see, e.g., [44]), the off-diagonal blocks of each summand $B = A - \mu_i I$ satisfy $\text{rank } B^{-1}|_{\text{off}} = \text{rank } B|_{\text{off}}$. Noting that $\text{rank } B|_{\text{off}} = b$, because B has bandwidth b , this completes the proof for simple poles. The result extends to non-simple poles by the semi-continuity of the rank function. \square

3.2.2 Singular value decay of off-diagonal blocks

The results of Lemma 4 and Lemma 5 allow us to establish exponential decay for the singular values of the off-diagonal blocks in $\Pi_{<0}(A)$ or, equivalently, in $\text{sign}(A)$ for any symmetric banded matrix A . By rescaling A , we may assume without loss of generality that its spectrum is contained in $[-R, -1] \cup [1, R]$. We let $\sigma_i(\cdot)$ denote the i th largest singular value of a matrix.

Theorem 6. *Consider a symmetric b -banded matrix $A \in \mathbb{R}^{n \times n}$ with the eigenvalues contained in $[-R, -1] \cup [1, R]$, and $m \in \mathbb{N}$. Letting $\text{gap} = 1/R$, the singular values of any off-diagonal block $\Pi_{<0}(A)|_{\text{off}}$ satisfy*

$$\sigma_{2mb+1}(\Pi_{<0}(A)|_{\text{off}}) \leq 2 \exp\left(-\frac{\pi^2 m}{4 \log(4/\sqrt[4]{\text{gap}} + 2)}\right).$$

Proof. Let s_m denote the solution of the min-max problem (9). Because $s_m(A)|_{\text{off}}$ has rank at most $2mb$ by Lemma 5, and the best rank- i approximation error is governed by the $(i + 1)$ th largest singular value, it follows from (12) that

$$\begin{aligned} \sigma_{2mb+1}(\text{sign}(A)|_{\text{off}}) &\leq \|\text{sign}(A) - s_m(A)\|_2 \leq \max_{x \in [-R, -1] \cup [1, R]} |\text{sign}(x) - s_m(x)| \\ &\leq 4 \exp\left(-\frac{\pi^2 m}{4 \log(4/\sqrt[4]{\text{gap}} + 2)}\right). \end{aligned}$$

The statement thus follows from the relation $\Pi_{<0}(A)|_{\text{off}} = -\frac{1}{2} \text{sign}(A)|_{\text{off}}$. \square

3.2.3 Memory requirements with respect to gap

Theorem 6 allows us to study the memory required to approximate $\Pi_{<0}(A)$ in the HODLR format to a prescribed accuracy. For this purpose, let $\Pi^{\mathcal{H}}$ denote the best approximation in the Frobenius norm of $\Pi_{<0}(A)$ in the HODLR format with all off-diagonal ranks bounded by $2mb$. Necessarily, the diagonal blocks of $\Pi^{\mathcal{H}}$ and $\Pi_{<0}(A)$ are the same. For an off-diagonal block of size k , Theorem 6 implies

$$\begin{aligned} \|\Pi_{<0}(A)|_{\text{off}} - \Pi^{\mathcal{H}}|_{\text{off}}\|_F^2 &= \sum_{i=2mb+1}^k \sigma_i(\Pi_{<0}(A)|_{\text{off}})^2 \leq \sum_{j=m}^{\lceil k/2b \rceil - m} 2b \sigma_{2jb+1}(\Pi_{<0}(A)|_{\text{off}})^2 \\ &\leq 8b \sum_{j=m}^{\lceil k/2b \rceil - m} \tau^{2j} \leq \frac{8b}{1 - \tau^2} \tau^{2m}, \end{aligned}$$

with $\tau = \exp\left(-\frac{\pi^2}{4\log(4/\sqrt[4]{\text{gap}+2})}\right)$. Taking into account the total number of off-diagonal blocks, we arrive at

$$\|\Pi_{<0}(A) - \Pi^{\mathcal{H}}\|_F^2 \leq \frac{16b}{1-\tau^2}(n/n_{\min} - 1)\tau^{2m}$$

Thus, the value of m needed to attain $\|\Pi_{<0}(A) - \Pi^{\mathcal{H}}\|_F \leq \delta$ for a desired accuracy $\delta > 0$ satisfies $m = \mathcal{O}(|\log \text{gap}| \cdot \log(bn\delta^{-1}|\log \text{gap}|))$. The corresponding approximation $\Pi^{\mathcal{H}}$ requires

$$\mathcal{O}\left(|\log \text{gap}| \cdot \log(bn\delta^{-1}|\log \text{gap}|)bn \log n\right) \quad (13)$$

memory. Up to a double logarithmic factor, this shows that the memory depends logarithmically on the spectral gap.

3.2.4 Comparison to approximate sparsity

We now compare (13) with known results for approximate sparsity. Assuming we are in the setting of Theorem 6, it is shown in [10] that the off-diagonal entries of $\Pi_{<0}(A)$ satisfy

$$|(\Pi_{<0}(A))_{ij}| \leq C e^{-\alpha|i-j|}, \quad \alpha = \frac{1}{2b} \log\left(\frac{1+\text{gap}}{1-\text{gap}}\right),$$

for some constant $C > 0$ depending only on R . Let $\Pi^{(m)}$ denote the best approximation in the Frobenius norm to $\Pi_{<0}(A)$ by a matrix of bandwidth m . Following [10, Theorem 7.7], we obtain

$$\|\Pi_{<0}(A) - \Pi^{(m)}\|_F \leq \frac{C}{\sqrt{\alpha}} \sqrt{n} e^{-\alpha m}.$$

Choosing a value of m that satisfies $m = \mathcal{O}(b\text{gap}^{-1} \log(Cbn\delta^{-1}\text{gap}^{-1}))$ thus ensures an accuracy of $\delta > 0$, where we used $\alpha \approx \text{gap}/b$. Since the storage of $\Pi^{(m)}$ requires $\mathcal{O}(mn)$ memory, we arrive at

$$\mathcal{O}\left(\frac{1}{\text{gap}} \log(Cbn\delta^{-1}\text{gap}^{-1})bn\right) \quad (14)$$

memory. In contrast to the logarithmic dependence in (13), the spectral gap now enters the asymptotic complexity inversely proportional. On the other hand, (13) features a factor $\log n$ that is not present in (14). For most situations of practical interest, we expect that the much milder dependence on the gap far outweighs this additional factor. In summary, the comparison between (13) and (14) provides strong theoretical justification for favoring the HODLR format over approximate sparsity.

4 QR-based first iteration of QDWH

The first QR-based iteration of the QDWH algorithm requires computing the QR decomposition

$$\begin{bmatrix} cA \\ I \end{bmatrix} = \begin{bmatrix} Q_1 \\ Q_2 \end{bmatrix} R \quad (15)$$

for some scalar $c > 0$. Without loss of generality, we suppose that $c = 1$. In this section, we develop an algorithm that requires $\mathcal{O}(b^2n)$ operations for performing this decomposition when A is a b -banded matrix. In particular, our algorithm directly computes Q_1 and Q_2 in the HODLR format. Since it is significantly simpler, we first discuss the case of a tridiagonal matrix A before treating the case of general b .

It is interesting to note that the need for computing a QR decomposition of the form (15) also arises in the solution of ill-posed inverse problems with Tikhonov regularization; see, e.g., [14]. However, when solving ill-posed problems, usually only the computation of the upper-triangular factor R is required, while the QDWH algorithm requires the computation of the orthogonal factor.

4.1 QR decomposition of $\begin{bmatrix} A \\ I \end{bmatrix}$ for tridiagonal A

For the case of a bidiagonal matrix A , Eldén [19] proposed a fast algorithm for reducing a matrix $\begin{bmatrix} A \\ I \end{bmatrix}$ to upper triangular form. In the following, we propose a modification of Eldén's algorithm suitable for tridiagonal A .

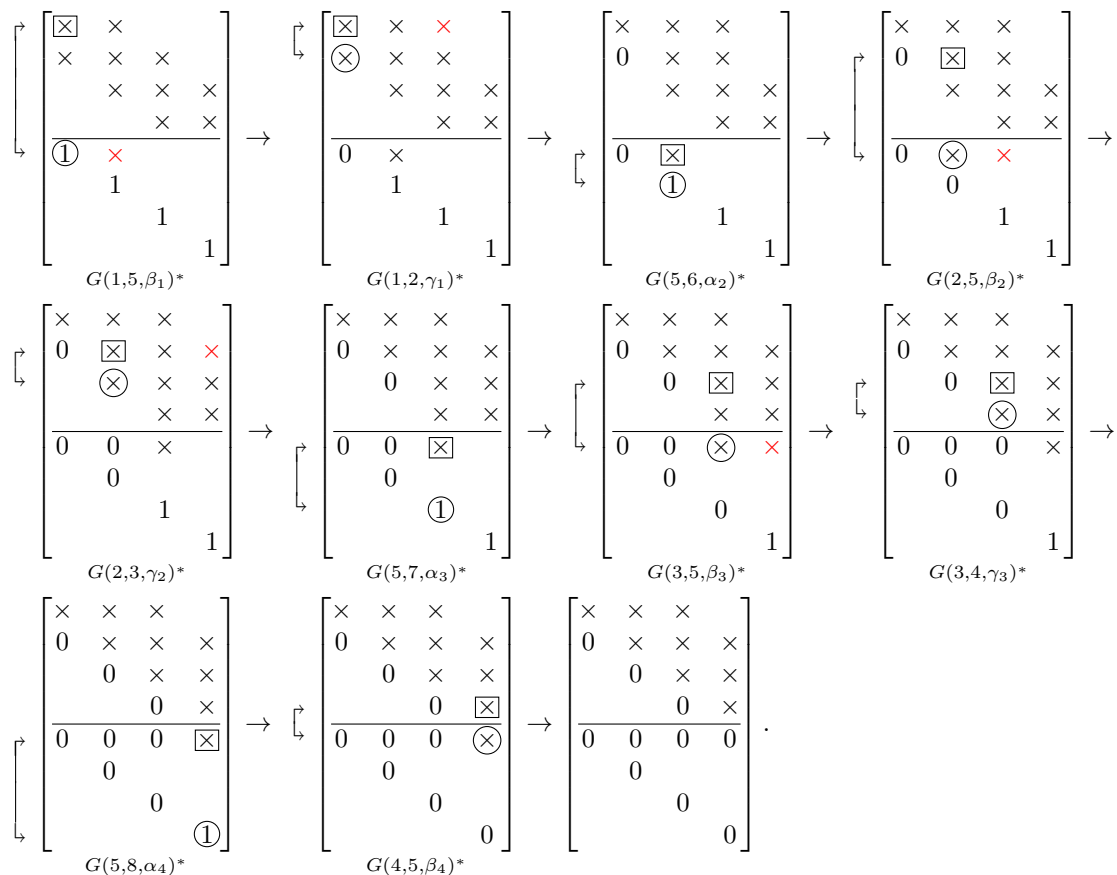


Figure 2: Fast QR decomposition of $\begin{bmatrix} A \\ I \end{bmatrix}$ for tridiagonal A and $n = 4$. In each step, a Givens rotation is applied to the rows denoted by the arrows. Crosses denote generically nonzero elements, boxed/circled crosses are used to define Givens rotations, while red crosses denote the fill-in during the current operation.

Our proposed algorithm is probably best understood from the illustration in Figure 2 for $n = 4$. In the i th step of the algorithm, all subdiagonal elements in the i th column of $\begin{bmatrix} A \\ I \end{bmatrix}$ are annihilated by performing Givens rotations either with the diagonal element, or with the element

$(n + 1, i)$. By carefully choosing the order of annihilation, only one new nonzero subdiagonal element is created in column $i + 1$. The detailed pseudocode of this procedure is provided in Algorithm 1. We use $G(i, j, \alpha)$ to denote a Givens rotation of angle α that is applied to rows/columns i and j .

Algorithm 1 Fast QR decomposition (15) for tridiagonal A

Input: Tridiagonal matrix A .

Output: Factors Q, R of a QR decomposition of $\begin{bmatrix} A \\ I \end{bmatrix}$.

```

1:  $Q \leftarrow I_{2n}, R \leftarrow \begin{bmatrix} A \\ I \end{bmatrix}$ .
2: Construct  $G(1, n + 1, \beta_1)$  to annihilate  $R(n + 1, 1)$ .
3: Update  $R \leftarrow G(1, n + 1, \beta_1)^* R$  and  $Q \leftarrow QG(1, n + 1, \beta_1)$ 
4: Construct  $G(1, 2, \gamma_1)$  to annihilate  $R(2, 1)$ .
5: Update  $R \leftarrow G(1, 2, \gamma_1)^* R$  and  $Q \leftarrow QG(1, 2, \gamma_1)$ .
6: for  $i = 2, \dots, n$  do
7:   Construct  $G(n + 1, n + i, \alpha_i)$  to annihilate  $R(n + i, i)$ .
8:   Update  $R \leftarrow G(n + 1, n + i, \alpha_i)^* R$  and  $Q \leftarrow QG(n + 1, n + i, \alpha_i)$ .
9:   Construct  $G(i, n + 1, \beta_i)$  to annihilate  $R(n + 1, i)$ .
10:  Update  $R \leftarrow G(i, n + 1, \beta_i)^* R$  and  $Q \leftarrow QG(i, n + 1, \beta_i)$ .
11:  if  $i < n$  then
12:    Construct  $G(i, i + 1, \gamma_i)$  to annihilate  $R(i + 1, i)$ .
13:    Update  $R \leftarrow G(i, i + 1, \gamma_i)^* R$  and  $Q \leftarrow QG(i, i + 1, \gamma_i)$ .
14:  end if
15: end for

```

Algorithm 1 performs $3n - 2$ Givens rotations in total. By exploiting its sparsity in a straightforward manner, only $\mathcal{O}(n)$ operations and memory are required to compute the upper triangular factor R . The situation is more complicated for the orthogonal factor. Since Q is dense, it would require $\mathcal{O}(n^2)$ operations and memory to form Q using Algorithm 1. In the following section, we explain how the low-rank structure of Q can be exploited to reduce this cost to $\mathcal{O}(n)$ as well.

4.1.1 Ranks of off-diagonal blocks and fast computation of orthogonal factor

For our purposes, it suffices to compute the first n columns of the $2n \times 2n$ matrix Q , that is, the $n \times n$ matrices $Q_1 = Q(1 : n, 1 : n)$ and $Q_2 = Q(n + 1 : 2n, 1 : n)$. The order of Givens rotations in Algorithm 1 implies that Q_1 is an upper Hessenberg matrix while Q_2 is an upper triangular matrix. The following theorem shows that all off-diagonal blocks of Q_1, Q_2 have rank at most two.

Theorem 7. *For the orthogonal factor Q returned by Algorithm 1, it holds that the matrices $Q(1 : k, k + 1 : n)$ and $Q(n + 1 : n + k, k + 1 : n)$ have rank at most two for all $1 \leq k < n$.*

Proof. We only prove the result for $Q(1 : k, k + 1 : n)$; the proof for $Q(n + 1 : n + k, k + 1 : n)$ is analogous.

During steps $1, \dots, k - 1$ of Algorithm 1, $Q(1 : k, k + 1 : n)$ is not modified and remains zero. In step k of Algorithm 1, column $k + 1$ of Q is modified, while $Q(1 : k, k + 2 : n)$ remains zero. After step k has been completed, let us set

$$\mathcal{U} := \text{span}\{Q(1 : k, k + 1), Q(1 : k, n + 1)\} \subset \mathbb{R}^k. \quad (16)$$

By construction, $\text{span } Q(1 : k, k + 1 : n) \subset \mathcal{U}$. In the following, we show by induction that this relation holds for all subsequent steps of Algorithm 1. Suppose that $\text{span } Q(1 : k, k + 1 : n) \subset \mathcal{U}$

holds after i steps for some i with $k \leq i \leq n-1$. In step $i+1$, the following operations are performed:

1. $G(n+1, n+i+1, \alpha_{i+1})$ is applied to columns $n+1$ and $n+i+1$ of Q . Because $Q(1:k, n+i+1)$ is zero before applying the rotation, this simply effects a rescaling of column $n+1$ and thus $Q(1:k, n+1) \in \mathcal{U}$ remains true.
2. $G(i+1, n+1, \beta_{i+1})$ is applied to columns $i+1$ and $n+1$ of Q , which preserves $\text{span } Q(1:k, k+1:n) \subset \mathcal{U}$.
3. If $i < n$, $G(i+1, i+2, \gamma_{i+1})$ is applied to columns $i+1$ and $i+2$ of Q , which again preserves $\text{span } Q(1:k, k+1:n) \subset \mathcal{U}$.

After completion of the algorithm, the column span of $Q(1:k, k+1:n)$ is thus contained in a subspace of dimension at most two. This proves the statement of the theorem. \square

Remark 8. *The proof of Theorem 7 can be turned into a procedure for directly computing low-rank representations for the off-diagonal blocks of Q_1, Q_2 in the HODLR format. Due to the structure of Q_1 and Q_2 , all lower off-diagonal blocks have ranks 1 and 0 respectively, and the computation of their low-rank representations is straightforward. In the following, we therefore only discuss the computation of a low-rank representation for an upper off-diagonal $p \times s$ block $Q_1|_{\text{off}} = U_1 V_1^*$ with $U_1 \in \mathbb{R}^{p \times 2}$, $V_1 \in \mathbb{R}^{s \times 2}$.*

Let $r+1$ and $k+1$ denote the row and column in Q_1 that correspond to the first row and column of $Q_1|_{\text{off}}$, respectively. The construction of $Q_1|_{\text{off}}$ begins in step k of Algorithm 1, because $Q_1|_{\text{off}}$ is zero before step k . During step k only the first column of $Q_1|_{\text{off}}$ is affected by $G(k, k+1, \gamma_k)$; it becomes a scalar multiple of $Q_1(r+1:r+p, k)$.

After step k of Algorithm 1 is completed, we set $U_1 = [Q_1|_{\text{off}}(:, 1), Q(r+1:r+p, n+1)]$, as in (16). The matrix V_1 stores the coefficients in the basis U_1 of the columns in $Q_1|_{\text{off}}$. Initially, $V_1 = [e_1, \mathbf{0}]$ with the first unit vector e_1 . As we also need to update the basis coefficients of $Q(r+1:r+p, n+1)$, we actually consider the augmented matrix $V^* = \begin{bmatrix} V_1^* & 0 \\ & 1 \end{bmatrix}$. In all subsequent steps of Algorithm 1, we only apply Givens rotations to the corresponding columns of V^* . Note that the last column of V^* is only rescaled, as it is always combined with a zero column. After completing step $k+s$ of Algorithm 1, $Q_1|_{\text{off}}$ remains unchanged and we extract the factor V_1 from the first s columns of V .

Using the described procedure, the overall complexity to compute a low rank representation of $Q_1|_{\text{off}}$ is $\mathcal{O}(\max\{p, s\})$. The off-diagonal blocks of Q_2 are treated analogously.

The QDWH algorithm makes use of the matrix product $Q_1 Q_2^*$, see (3b). Theorem 7, together with the upper Hessenberg/triangular structure, directly implies that the ranks of the lower and upper off-diagonal blocks of $Q_1 Q_2^*$ are bounded by three and two, respectively. In fact, the following theorem shows a slightly stronger result.

Theorem 9. *Every off-diagonal block of $Q_1 Q_2^*$ for the orthogonal factor returned by Algorithm 1 has rank at most 2.*

Proof. By Algorithm 1 and Theorem 7, the matrices Q_1 and Q_2 admit for any $1 \leq k < n$ a partitioning of the form

$$Q_1 = \left[\begin{array}{c|c} X_1 & U_1 V_1^* \\ \hline \sigma e_1 e_k^* & X_2 \end{array} \right], \quad Q_2 = \left[\begin{array}{c|c} Y_1 & U_2 V_2^* \\ \hline \mathbf{0} & Y_2 \end{array} \right],$$

where $X_1 \in \mathbb{R}^{k \times k}$, $X_2 \in \mathbb{R}^{n-k \times n-k}$ are upper Hessenberg, $Y_1 \in \mathbb{R}^{k \times k}$, $Y_2 \in \mathbb{R}^{n-k \times n-k}$ are upper triangular, $U_1, U_2 \in \mathbb{R}^{k \times 2}$, $V_1, V_2 \in \mathbb{R}^{n-k \times 2}$, $\sigma \in \mathbb{R}$, and e_1, e_k denote unit vectors of appropriate

lengths. The upper off-diagonal block of $Q_1 Q_2^*$ equals to a product of rank-2 matrices

$$(Q_1 Q_2^*)(1 : k, k + 1 : n) = X_1 \cdot \mathbf{0} + U_1 \underbrace{V_1^* Y_2^*}_{\tilde{V}_1^*} = U_1 \tilde{V}_1^*.$$

Moreover, the lower off-diagonal block amounts to a sum of a rank-1 and a rank-2 matrix

$$\begin{aligned} (Q_1 Q_2^*)(k + 1 : n, 1 : k) &= \sigma e_1 e_k^* Y_1^* + \underbrace{X_2 V_2 U_2^*}_{\tilde{V}_2} \\ &= \sigma e_1 Y_1(:, k)^* + \tilde{V}_2 U_2^*. \end{aligned}$$

If $\sigma = 0$, the statement holds. Otherwise, we first show that the vectors $Y_1(:, k)$ and $U_2(:, 1)$ are collinear. Let us recall that the vectors $Y_1(:, k)$, $U_2(:, 1)$ coincide with the vectors $Q(n + 1 : n + k + 1, k)$, $Q(n + 1 : n + k + 1, k + 1)$ computed during step k of Algorithm 1. As $Q(n + 1 : n + k + 1, k)$ and $Q(n + 1 : n + k + 1, k + 1)$ are collinear after performing step k , the same holds for $Y_1(:, k)$, $U_2(:, 1)$, and $Y_1(:, k) = \eta U_2(:, 1)$ for some $\eta \in \mathbb{R}$. Hence, we obtain

$$(Q_1 Q_2^*)(k + 1 : n, 1 : k) = \sigma \eta e_1 U_2(:, 1)^* + \tilde{V}_2 U_2^* = \hat{V}_2 U_2^*,$$

which completes the proof. \square

From Theorem 9 and the recurrence (3) it follows that the first iterate of the QDWH algorithm can be exactly represented in the HODLR format with off-diagonal ranks at most 3.

4.2 QR decomposition of $\begin{bmatrix} A \\ I \end{bmatrix}$ for banded A

In this section, we discuss the QR decomposition of $\begin{bmatrix} A \\ I \end{bmatrix}$ for a banded symmetric matrix A with bandwidth $b > 1$. Let us first note that Eldén [20] proposed a fast algorithm for reducing a matrix $\begin{bmatrix} A \\ L \end{bmatrix}$ with an upper triangular banded matrix L . Eldén's algorithm does not cover the fast computation of the orthogonal factor and requires the application of $(2b + 1)n + nb - \frac{5}{2}b^2 - \frac{3}{2}b$ Givens rotations. In the following, we propose a different algorithm that only requires $(2b + 1)n - b^2 - b$ Givens rotations.

Figure 10 illustrates the idea of our algorithm for $n = 6$ and $b = 3$. In the i th step of the algorithm, the subdiagonal elements in the i th column of $\begin{bmatrix} A \\ I \end{bmatrix}$ are annihilated as follows. A first group of Givens rotations $(\alpha_{i,j})$ annihilates all elements in row $n + i$, which consists of the diagonal element of I and fill-in from the previous step. Then a Givens rotation (β_i) annihilates the element $(n + 1, i)$. Finally, a second group of Givens rotations $(\gamma_{i,j})$ annihilates all subdiagonal elements of A . The detailed procedure is given in Algorithm 2.

4.2.1 Ranks of off-diagonal blocks and fast computation of orthogonal factor

Due to the order of annihilation in Algorithm 2, it follows that $Q_1 = Q(1 : n, 1 : n)$ is a b -Hessenberg matrix (that is, the matrix is zero below the b th subdiagonal) while $Q_2 = Q(n + 1 : 2n, 1 : n)$ is an upper triangular matrix. The following result and its proof yield an $\mathcal{O}(b^2 n)$ algorithm for computing Q_1 and Q_2 , analogous to Theorem 7 and Remark 8.

Theorem 10. *For the orthogonal factor Q returned by Algorithm 2, it holds that the matrices $Q(1 : k, k + 1 : n)$ and $Q(n + 1 : n + k, k + 1 : n)$ have rank at most $2b$ for all $1 \leq k < n$.*

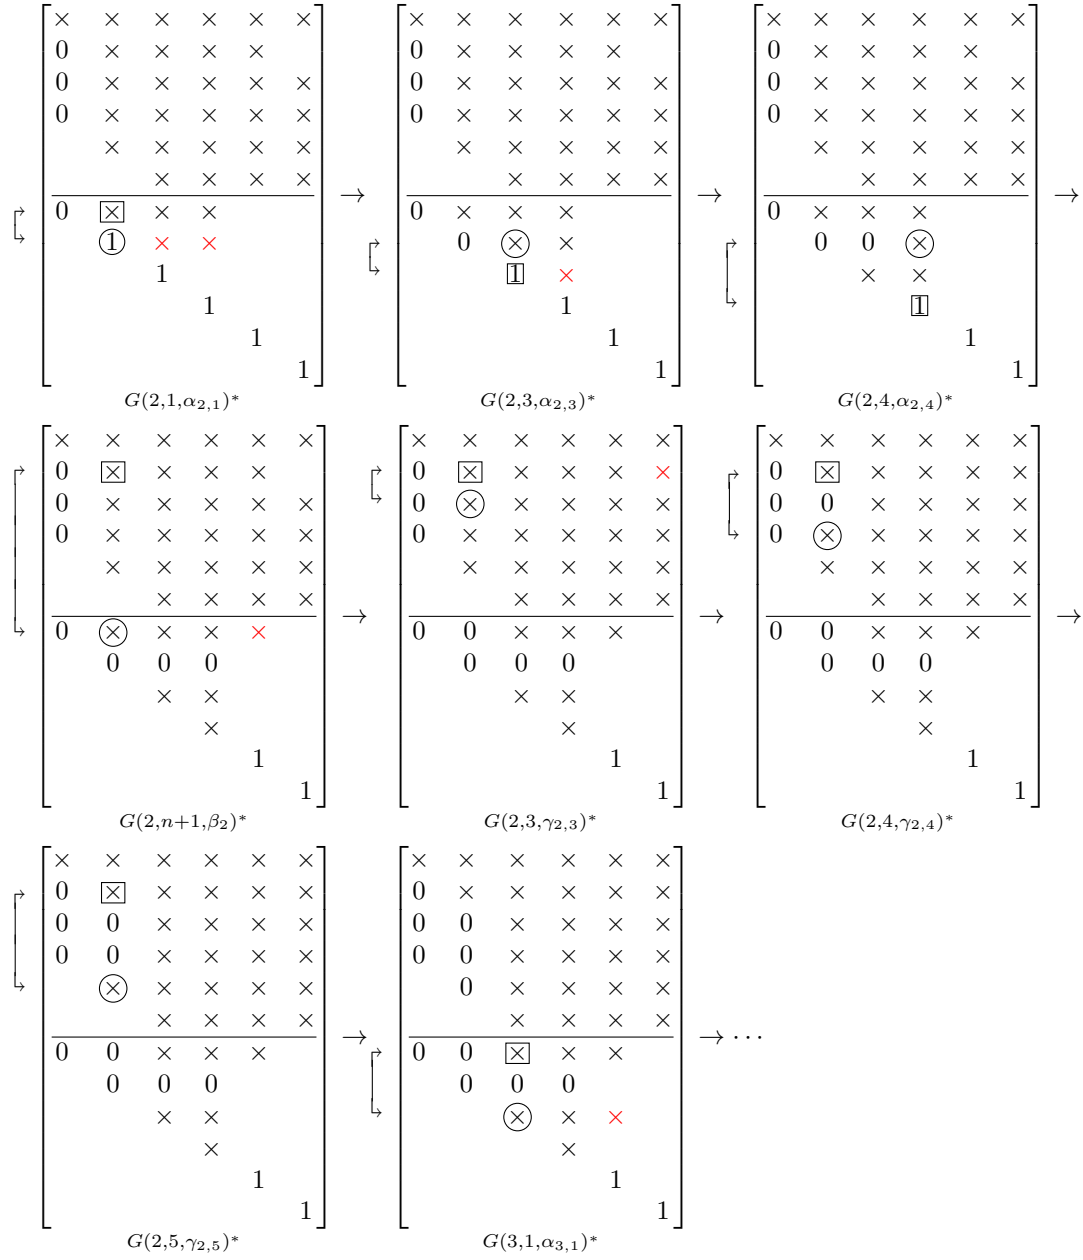


Figure 3: Second step of fast QR decomposition (Algorithm 2) of $\begin{bmatrix} A \\ I \end{bmatrix}$ for banded A with $n = 6$ and $b = 3$. In each step, a Givens rotation is applied to the rows denoted by the arrows. Crosses denote generically nonzero elements, boxed/circled crosses are used to define Givens rotations, while red crosses denote the fill-in during the current operation.

Proof. Again, we prove the result for $Q(1 : k, k + 1 : n)$ only. After k steps of Algorithm 2 have

Algorithm 2 Fast QR decomposition (15) for banded A

Input: Banded matrix A with bandwidth b .

Output: Factors Q, R of a QR decomposition of $\begin{bmatrix} A \\ I \end{bmatrix}$.

```

1:  $Q \leftarrow I_{2n}, R \leftarrow \begin{bmatrix} A \\ I \end{bmatrix}$ .
2: Construct  $G(1, n+1, \beta_1)$  to annihilate  $R(n+1, 1)$ .
3: Update  $R \leftarrow G(1, n+1, \beta_1)^* R$  and  $Q \leftarrow QG(1, n+1, \beta_1)$ 
4: for  $j = 2, \dots, b+1$  do
5:   Construct  $G(1, j, \gamma_{1,j})$  to annihilate  $R(j, 1)$ .
6:   Update  $R \leftarrow G(1, j, \gamma_{1,j})^* R$  and  $Q \leftarrow QG(1, j, \gamma_{1,j})$ .
7: end for
8: for  $i = 2, \dots, n$  do
9:   Construct  $G(n+1, n+i, \alpha_{i,i})$  to annihilate  $R(n+i, i)$ .
10:  Update  $R \leftarrow G(n+1, n+i, \alpha_{i,i})^* R$  and  $Q \leftarrow QG(n+1, n+i, \alpha_{i,i})$ .
11:  for  $j = i+1, \dots, \min\{n, b+i-1\}$  do
12:    Construct  $G(n+i, n+j, \alpha_{i,j})$  to annihilate  $R(n+i, j)$ .
13:    Update  $R \leftarrow G(n+i, n+j, \alpha_{i,j})^* R$  and  $Q \leftarrow QG(n+i, n+j, \alpha_{i,j})$ .
14:  end for
15:  Construct  $G(i, n+1, \beta_i)$  to annihilate  $R(n+1, i)$ .
16:  Update  $R \leftarrow G(i, n+1, \beta_i)^* R$  and  $Q \leftarrow QG(i, n+1, \beta_i)$ .
17:  if  $i < n$  then
18:    for  $j = i+1, \dots, \min\{n, b+i\}$  do
19:      Construct  $G(i, j, \gamma_{i,j})$  to annihilate  $R(j, i)$ .
20:      Update  $R \leftarrow G(i, j, \gamma_{i,j})^* R$  and  $Q \leftarrow QG(i, j, \gamma_{i,j})$ .
21:    end for
22:  end if
23: end for

```

been performed, we define the subspace

$$\mathcal{U} := \text{span}\{Q(1:k, k+1), \dots, Q(1:k, k+b), Q(1:k, n+1), \\ Q(1:k, n+k+1), \dots, Q(1:k, n+k+b-1)\},$$

which is of dimension not larger than $2b$. At this point, the columns $Q(1:k, j)$ are zero for $j = k+b+1, \dots, n$ and $j = n+k+b, \dots, 2n$. Thus,

$$\text{span } Q(1:k, k+1:n+1) \subset \mathcal{U}, \quad \text{span } Q(1:k, n+k+1:2n) \subset \mathcal{U} \quad (17)$$

hold after k steps of Algorithm 2. We now show by induction that this relation holds for all subsequent steps.

Suppose that (17) holds after i steps with $k \leq i \leq n-1$. In step $i+1$, the following operations are performed by Algorithm 2:

1. $G(n+1, n+i+1, \alpha_{i+1, i+1})$ is applied to columns $n+1$ and $n+i+1$ of Q , which affects and preserves both inclusions in (17). Then $G(n+i+1, n+j, \alpha_{i+1, j})$ is applied to columns $n+i+1$ and $n+j$ of Q , for $j = i+2 : \min\{n, i+b\}$, hence $\text{span } Q(1:k, n+k+1:2n) \subset \mathcal{U}$ remains true.
2. $G(i+1, n+1, \beta_{i+1})$ is applied to columns $i+1$ and $n+1$ of Q , preserving $\text{span } Q(1:k, k+1:n+1) \subset \mathcal{U}$.

3. If $i + 1 < n$, $G(i + 1, j, \gamma_{i+1, j})$ is applied to columns $i + 1$ and j of Q , for $j = i + 2 : \min\{n, i + b + 1\}$, which retains $\text{span } Q(1 : k, k + 1 : n + 1) \subset \mathcal{U}$.

Therefore (17) holds after Algorithm 2 has been completed, which completes the proof of the theorem. \square

The following result is an extension of Theorem 11 from the tridiagonal to the banded case. Its proof is very similar and therefore omitted.

Theorem 11. *Every off-diagonal block of $Q_1 Q_2^*$ for the orthogonal factor returned by Algorithm 2 has rank at most $2b$.*

5 hQDWH algorithm

Algorithm 3, summarizes the hQDWH algorithm proposed in this paper.

Algorithm 3 hQDWH algorithm

Input: Symmetric banded matrix A with bandwidth $b \geq 1$, minimal block-size $n_{\min} \geq 2$, truncation tolerance $\epsilon > 0$, stopping tolerance $\delta > 0$.

Output: Approximation P in HODLR format to spectral projector $\Pi_{<0}(A)$.

- 1: Choose initial parameters α, l_0 of QDWH according to (18).
 - 2: $X_0 = A/\alpha$.
 - 3: $k = 0$.
 - 4: **while** $|1 - l_k| > \delta$ **do**
 - 5: Compute a_k, b_k, c_k according to the recurrence (2).
 - 6: **if** $k = 0$ **then**
 - 7: Apply $\begin{cases} \text{Algorithm 1,} & \text{for } b = 1 \\ \text{Algorithm 2,} & \text{for } b > 1 \end{cases}$ to $\begin{bmatrix} \sqrt{c_0} X_0 \\ I \end{bmatrix}$ and store resulting array G of Givens rotations.
 - 8: Compute Q_1 and Q_2 from G in HODLR format; see Remark 8.
 - 9: $X_1 = \frac{b_0}{c_0} *_{\mathcal{H}} X_0 +_{\mathcal{H}} \frac{1}{\sqrt{c_0}} \left(a_0 - \frac{b_0}{c_0} \right) *_{\mathcal{H}} Q_1 *_{\mathcal{H}} Q_2^*$.
 - 10: **else**
 - 11: $W_k = \mathcal{H}\text{-Cholesky}(I +_{\mathcal{H}} c_k *_{\mathcal{H}} X_k^* *_{\mathcal{H}} X_k)$.
 - 12: Solve upper-triangular system $Y_k W_k = X_k$ in HODLR format.
 - 13: Solve lower-triangular system $V_k W_k^* = Y_k$ in HODLR format.
 - 14: $X_{k+1} = \frac{b_k}{c_k} *_{\mathcal{H}} X_k +_{\mathcal{H}} \left(a_k - \frac{b_k}{c_k} \right) *_{\mathcal{H}} V_k$.
 - 15: **end if**
 - 16: $k = k + 1$.
 - 17: $l_k = l_{k-1} (a_{k-1} + b_{k-1} l_{k-1}^2) / (1 + c_{k-1} l_{k-1}^2)$.
 - 18: **end while**
 - 19: $U = X_k$.
 - 20: $P = \frac{1}{2} *_{\mathcal{H}} (I -_{\mathcal{H}} U)$.
-

In the following, we comment on various implementation details of Algorithm 3.

line 1 As proposed in [39], the parameters $\alpha \gtrsim \|A\|_2$ and $l_0 \lesssim \sigma_{\min}(X_0)$ needed to start the QDWH algorithm are estimated as

$$\alpha = \text{normest}(A), \quad l_0 = \|A/\alpha\|_1 / (\sqrt{n} \cdot \text{condest}(A/\alpha)), \quad (18)$$

where `normest` and `condest` denote the MATLAB functions for estimating the matrix 2-norm using the power method and the 1-norm condition number using [32], respectively. Both functions exploit that A is sparse and require $\mathcal{O}(bn)$ and $\mathcal{O}(b^2n)$ operations, respectively.

lines 7– 9 This part of the algorithm deals with the implementation of the first QR-based iterate (3). The generation of Givens rotations by Algorithms 1 and 2 for reducing $\begin{bmatrix} \sqrt{c_0} X_0 \\ I \end{bmatrix}$ to triangular form has been implemented in a C function, making use of the LAPACK routine DLARTG. The function is called via a MEX interface and returns an array G containing the cosines and sines of all rotations. This array is then used in 8 to generate Q_1 and Q_2 in the HODLR format, whose precise form is defined by the input parameter n_{\min} .

lines 11– 14 The computation of the k th iterate X_k , $k > 1$, involves the Cholesky decomposition, addition, and the solution of triangular linear systems in the HODLR format. Existing techniques for HODLR matrices have been used for this purpose, see Section 3, and repeated recompression with the absolute truncation tolerance ϵ is applied.

Remark 12. *Algorithm 3 extends in a straightforward way to the more general hierarchical matrix format from Section 3.1.2. The only major difference is the need for converting the matrices after line 9 from the HODLR to the hierarchical matrix format. This extension of Algorithm 3 was used in Example 3.*

Assuming that all ranks in the off-diagonal blocks are bounded by $k \geq b$, Algorithm 3 requires $\mathcal{O}(kn \log n)$ memory and $\mathcal{O}(k^2 n \log^2 n)$ operations.

6 Numerical experiments

In this section, we demonstrate the performance of our preliminary MATLAB implementation of the hQDWH algorithm. All computations were performed in MATLAB version 2014a on an Intel Xeon CPU with 3.07GHz, 4096 KByte of level 2 cache and 192 GByte of RAM. To be able to draw a fair comparison, all experiments were performed on a single core.

To measure the accuracy of the QDWH algorithm, we use the functions e_{id}^Q , e_{trace}^Q , e_{SP}^Q defined in (6). The error measures e_{id}^H , e_{trace}^H , e_{SP}^H for the hQDWH algorithm are defined analogously. In all experiments, we used the tolerance $\delta = 10^{-15}$ for stopping the QDWH/hQDWH algorithms. Unless stated otherwise, the truncation tolerance for recompression in the HODLR format is set to $\epsilon = 10^{-10}$; the minimal block-size is set to $n_{\min} = 250$ for tridiagonal matrices and $n_{\min} = 500$ for banded matrices.

The performance of the algorithm is tested on various types of matrices, including synthetic examples as well as examples from widely used sparse matrix collections.

6.1 Construction of synthetic test matrices

Given a prescribed set of eigenvalues $\lambda_1, \dots, \lambda_n$ and a bandwidth b , we construct a symmetric b -banded matrix by an orthogonal similarity transformation of $A = \text{diag}(\lambda_1, \dots, \lambda_n)$. For this purpose, we perform the following operation for $i = n, n-1, \dots, 2$:

First, a Givens rotations $G(i-1, i, \alpha_i)$ is created by annihilating the second component of the vector $\begin{bmatrix} a_{ii} \\ 1 \end{bmatrix}$. The update $A \leftarrow G(i-1, i, \alpha_i) * A G(i-1, i, \alpha_i)$ introduces nonzero off-diagonal elements in A . For $i = n, \dots, n-b+1$, this fill-in stays within the b bands. For $i \leq n-b$, two undesired nonzero elements are created in row $i-1$ and column $i-1$ outside the b bands. These

nonzero elements are immediately chased off to the bottom right corner by applying $n - b - i + 1$ Givens rotations, akin to Schwarz band reduction [42].

When the procedure is completed, the b bands of A are fully populated.

In all examples below, we choose the eigenvalues to be uniformly distributed in $[-1, -\text{gap}] \cup [\text{gap}, 1]$. Our results indicate that the performance of our algorithm is robust with respect to the choice of eigenvalue distribution. In particular, the timings stay almost the same when choosing a distribution geometrically graded towards the spectral gap.

6.2 Results for tridiagonal matrices

Example 13 (Accuracy versus gap). First we investigate the behavior of the errors for hQDWH and QDWH with respect to the spectral gap. Using the construction from Section 6.1, we consider 10000×10000 tridiagonal matrices with eigenvalues in $[-1, -\text{gap}] \cup [\text{gap}, 1]$, where gap varies from 10^{-15} to 10^{-1} . From Figure 4 (left), it can be seen that tiny spectral gaps do not have a significant influence on the distance from identity and the trace error for both algorithms. On the other hand, both $e_{\text{SP}}^{\mathcal{H}}$ and $e_{\text{SP}}^{\mathcal{Q}}$ are sensitive to a decreasing gap, which reflects the ill-conditioning of the spectral projector for small gaps.

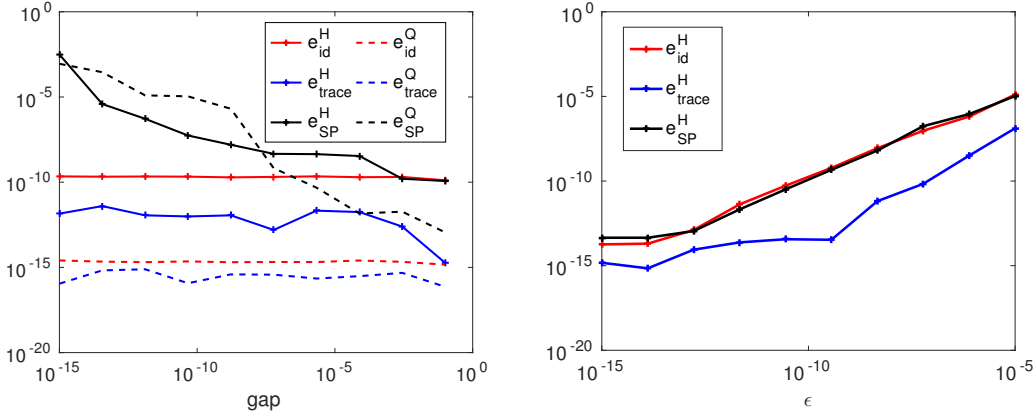


Figure 4: Left (Example 13): Comparison of accuracy for hQDWH and QDWH applied to tridiagonal matrices. Right (Example 14): Accuracy of hQDWH for different truncation tolerances.

Example 14 (Accuracy versus ϵ). We again consider a tridiagonal matrix $A \in \mathbb{R}^{10000 \times 10000}$, with the eigenvalues in $[-1, -10^{-4}] \cup [10^{-4}, 1]$. The truncation tolerance ϵ for recompression in the HODLR format is varied in the interval $[10^{-15}, 10^{-5}]$. Figure 4 shows the resulting errors in the hQDWH algorithm. As expected, the errors $e_{\text{id}}^{\mathcal{H}}$, $e_{\text{trace}}^{\mathcal{H}}$, $e_{\text{SP}}^{\mathcal{H}}$ increase as ϵ increases. Both $e_{\text{id}}^{\mathcal{H}}$ and $e_{\text{SP}}^{\mathcal{H}}$ grow linearly with respect to ϵ , while $e_{\text{trace}}^{\mathcal{H}}$ appears to be a little more robust.

Example 15 (Accuracy for examples from matrix collections). We tested the accuracy of the hQDWH algorithm for the matrices from applications also considered in [36]:

- Matrices from the BCSSTRUC1 set in the Harwell-Boeing Collection [16]. In these examples, a finite element discretization leads to a generalized eigenvalue problem $Kx = \lambda Mx$ with K, M symmetric positive definite. We consider the equivalent standard eigenvalue problem $L^{-1}KL^{-T}$, where L denotes the Cholesky factor of M , and shift the matrix such

that approximately half of its spectrum is negative. Finally, the matrix is reduced to a tridiagonal matrix using the Matlab function `hess`.

- Matrices from UF Sparse Matrix Collection [16]. We consider the symmetric Alemdar and Cannizzo matrices, as well as a matrix from the NASA set. Again, the matrices are shifted such that roughly half of their spectrum is negative and then reduced to tridiagonal form.

	matrix	n	$e_{\text{id}}^{\mathcal{H}}$	$e_{\text{trace}}^{\mathcal{H}}$	$e_{\text{SP}}^{\mathcal{H}}$	$\ \cdot\ _2$	gap
BCSSTRU1	bcsst08	1074	10^{-11}	10^{-13}	10^{-9}	$1.68 \cdot 10^7$	1.2615
	bcsst09	1083	10^{-10}	10^{-11}	10^{-8}	$4.29 \cdot 10^{14}$	$2.13 \cdot 10^{10}$
	bcsst11	1474	10^{-10}	10^{-12}	10^{-7}	$4.75 \cdot 10^9$	$1.58 \cdot 10^5$
	Cannizzo matrix	4098	10^{-10}	10^{-10}	10^{-7}	$3.07 \cdot 10^8$	0.728
	nasa4704	4704	10^{-10}	10^{-12}	10^{-9}	$2.07 \cdot 10^8$	20.182
	Alemdar matrix	6245	10^{-10}	10^{-11}	10^{-7}	69.5	0.0079

Table 4: Accuracy of hQDWH for tridiagonal matrices from Example 15.

Table 4 reveals that the hQDWH algorithm yields accurate approximations also for applications' matrices. More specifically, in all examples, $e_{\text{id}}^{\mathcal{H}}$ and $e_{\text{trace}}^{\mathcal{H}}$ obtain values of order of ϵ , while $e_{\text{SP}}^{\mathcal{H}}$ shows dependence on the relative spectral gap.

Example 16 (Breakeven point relative to eig). To compare the computational times of the hQDWH algorithm with `eig`, we consider tridiagonal matrices with eigenvalues contained in $[-1, -\text{gap}] \cup [\text{gap}, 1]$ for various gaps. Table 5 shows the resulting breakeven points, that is, the value of n such that hQDWH is faster than `eig` for matrices of size at least n . Not surprisingly, this breakeven point depends on the gap, as ranks are expected to increase as the gap decreases. However, even for $\text{gap} = 10^{-4}$, the hQDWH algorithm becomes faster than `eig` for matrices of moderate size ($n \geq 3250$) and the ranks of the off-diagonal blocks in the HODLR representation of the spectral projector remain reasonably small.

gap	breakeven point	max off-diagonal rank
10^{-1}	$n = 2250$	18
10^{-2}	$n = 2500$	28
10^{-3}	$n = 2750$	35
10^{-4}	$n = 3250$	37

Table 5: Breakeven point of hQDWH relative to `eig` for tridiagonal matrices. The last column shows the maximal off-diagonal rank in the output of hQDWH.

Example 17 (Performance versus n). In this example, we investigate the asymptotic behavior of the hQDWH algorithm, in terms of computational time and memory, for tridiagonal matrices with eigenvalues in $[-1, -10^{-6}] \cup [10^{-6}, 1]$. Figure 5 indicates that the expected $\mathcal{O}(n \log^2 n)$ computational time and $\mathcal{O}(n \log n)$ are nicely matched. The faster increase for smaller n is due to fact that the off-diagonal ranks first grow from 30 to 64 until they settle around 64 for sufficiently large n .

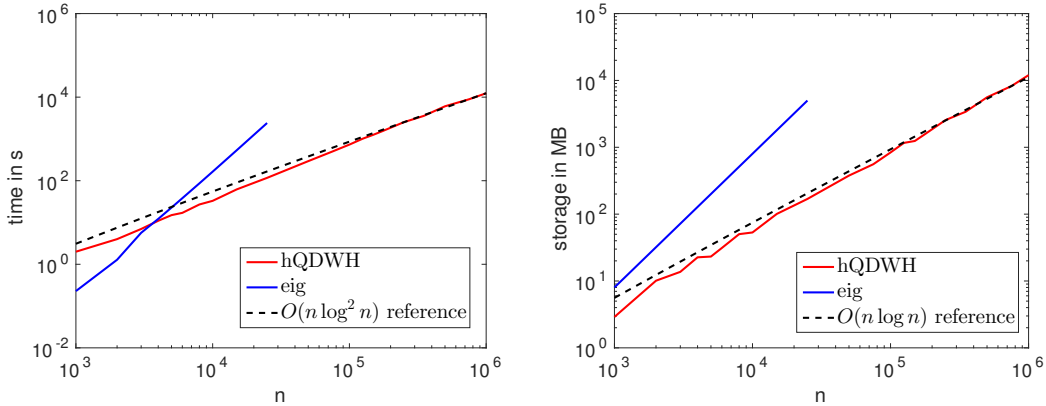


Figure 5: Example 17. Performance of hQDWH and *eig* applied to tridiagonal matrices with respect to n . Left: Computational time. Right: Memory requirements.

Example 18 (Performance for 1D Laplace). It is interesting to test the performance of the hQDWH algorithm for matrices for which the spectral gap decreases as n increases. The archetypical example is the (scaled) tridiagonal matrix from the central difference discretization of the 1D Laplace operator, with eigenvalues $\lambda_k = 2 - 2 \cos \frac{k\pi}{n+1}$ for $k = 1, \dots, n$. The matrix is shifted by 2, such that half of its spectrum is negative and the eigenvalues become equal to $\lambda_k = -2 \cos \frac{k\pi}{n+1}$. The spectral gap is given by $\text{gap} = 2 \sin \frac{\pi}{n+1} = \mathcal{O}(1/n^2)$. According to Theorem 6, the numerical ranks of the off-diagonal blocks depend logarithmically on the spectral gap. Thus, we expect that the hQDWH algorithm requires $\mathcal{O}(n \log^4 n)$ computational time and $\mathcal{O}(n \log^2 n)$ memory for this matrix. Figure 6 nicely confirms this expectation.

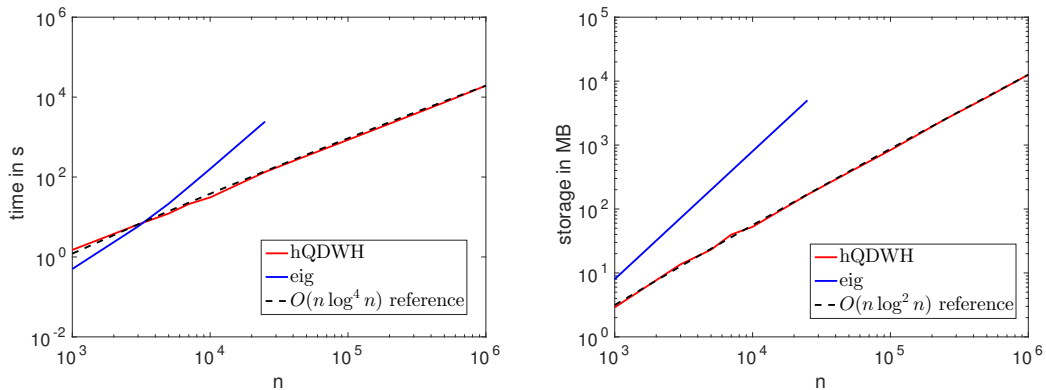


Figure 6: Example 18. hQDWH and *eig* for discretized 1D Laplace. Left: Computational time with respect to n . Right: Memory requirements with respect to n .

Example 19 (Performance versus n_{\min}). The choice of the minimal block size n_{\min} in the HODLR format influences the performance of hQDWH. We have investigated this dependence for $50\,000 \times 50\,000$ tridiagonal matrices with eigenvalues contained in $[-1, -\text{gap}] \cup [\text{gap}, 1]$, and

$\text{gap} \in \{10^{-1}, 10^{-4}, 10^{-6}\}$. Figure 7 indicates that the optimal value of n_{\min} increases for smaller gaps. However, the execution time is not overly sensitive to this choice; a value of n_{\min} between 200 and 500 leads to good performance.




Figure 7: Example 19. Computational time of hQDWH versus n_{\min} .

6.3 Results for banded matrices

Example 20 (Accuracy versus gap). Similarly to Example 13, we study the impact of the spectral gap on the accuracy of hQDWH and QDWH for banded matrices. Using once again the construction from Section 6.1, we consider 10000×10000 banded matrices with bandwidth 8 and eigenvalues in $[-1, -\text{gap}] \cup [\text{gap}, 1]$, where gap varies from 10^{-15} to 10^{-1} . The left plot of Figure 8 reconfirms the observations from Example 13.

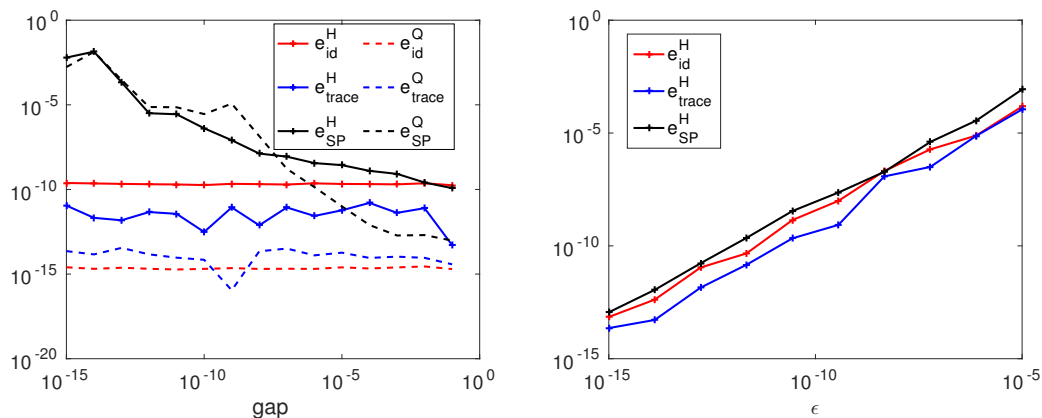


Figure 8: Left (Example 20): Comparison of accuracy for hQDWH and QDWH applied to banded matrices with bandwidth 8. Right (Example 21): Accuracy of hQDWH for different truncation tolerances.

Example 21 (Accuracy versus ϵ). We investigate the influence of the truncation tolerance ϵ on accuracy for an $10\,000 \times 10\,000$ banded matrix with bandwidth $b = 8$ and the eigenvalues contained in $[-1, -10^{-4}] \cup [10^{-4}, 1]$. The right plot of Figure 8 reconfirms the observations from Example 14.

Example 22 (Breakeven point relative to eig). Table 6 shows when hQDWH becomes faster than eig for $n \times n$ banded matrices with eigenvalues contained in $[-1, -\text{gap}] \cup [\text{gap}, 1]$ for $\text{gap} = 10^{-1}$ and 10^{-4} . Compared to Table 5, the breakeven point is lower for bandwidths $b = 2$ and $b = 4$ than for bandwidth 1. This is because eig needs to perform tridiagonal reduction when $b \geq 2$.

gap \ b	2	4	8	16
10^{-1}	$n = 1250$	$n = 1750$	$n = 2500$	$n = 5250$
10^{-4}	$n = 1750$	$n = 2500$	$n = 5000$	$n = 9500$

Table 6: Breakeven point of hQDWH relative to eig applied for banded matrices with various bandwidths and spectral gaps.

Example 23 (Performance versus n). We consider banded matrices with bandwidth 4 and with eigenvalues contained in $[-1, -10^{-1}] \cup [10^{-1}, 1]$. As in Example 17, Figure 9 confirms that the computational time of hQDWH scales like $\mathcal{O}(n \log^2 n)$ while memory scales like $\mathcal{O}(n \log n)$. Note that the maximal rank in the off-diagonal blocks is 66 for $n = 1\,000$, and 71 for $n = 500\,000$.

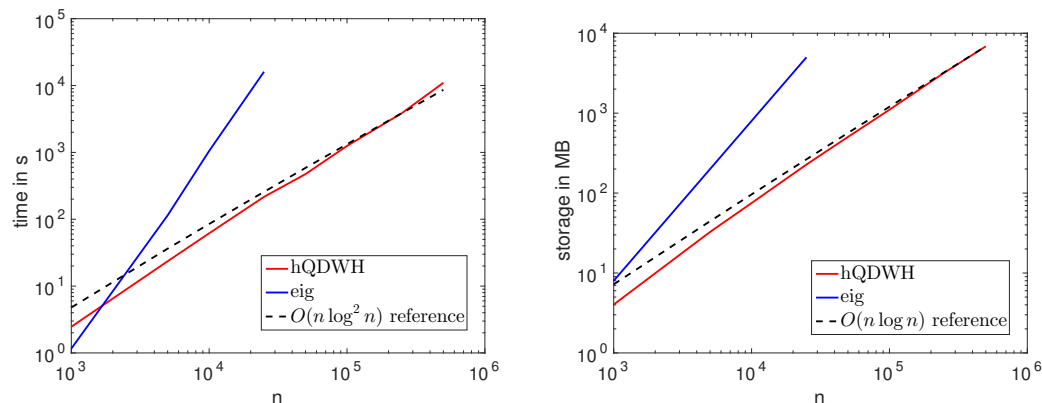


Figure 9: Example 23. Performance with respect to n of hQDWH and eig applied to banded matrices with bandwidth 4. Left: Computational time. Right: Memory requirements.

Example 24 (Performance versus b). To verify the influence of the matrix bandwidth on the performance of our algorithm, we consider $100\,000 \times 100\,000$ banded matrices with eigenvalues contained in $[-1, -10^{-6}] \cup [10^{-6}, 1]$. Figure 10 clearly demonstrates that computational time grows quadratically while memory grows linearly with respect to the bandwidth b .




Figure 10: Example 24. Performance with respect to bandwidth b of hQDWH applied to $100\,000 \times 100\,000$ banded matrices. Left: Computational time. Right: Memory requirements.

7 Conclusion

In this paper we have developed a fast algorithm for computing spectral projectors of large-scale symmetric banded matrices. For this purpose, we have tailored the ingredients of the QDWH algorithm, such that the overall algorithm has linear-polylogarithmic complexity. This allows us to compute highly accurate approximations to the spectral projector for very large sizes (up to $n = 1\,000\,000$ on a desktop computer) even when the involved spectral gap becomes small.

The choice of hierarchical low-rank matrix format is critical to the performance of our algorithm. Somewhat surprisingly, we have observed that the relatively simple HODLR format outperforms a more general \mathcal{H} -matrix format. We have not investigated the choice of a format with nested low-rank factors, such as HSS matrices. While such a nested format likely lowers asymptotic complexity, it presumably only pays off for larger values of n .

Acknowledgements. We are grateful to Jonas Ballani, Petar Sirković, and Michael Steinlechner for helpful discussions on this paper, as well as to Stefan Güttel for providing us insights into the approximation error results used in Section 3.2.1.

References

- [1] N. I. Akhiezer. *Elements of the Theory of Elliptic Functions*, volume 79 of *Translations of Mathematical Monographs*. American Mathematical Society, Providence, RI, 1990.
- [2] P. Arbenz. Divide and conquer algorithms for the bandsymmetric eigenvalue problem. *Parallel Comput.*, 18(10):1105–1128, 1992.
- [3] T. Auckenthaler, V. Blum, H.-J. Bungartz, T. Huckle, R. Johanni, L. Krämer, B. Lang, H. Lederer, and P. R. Willems. Parallel solution of partial symmetric eigenvalue problems from electronic structure calculations. *Parallel Computing*, 37(12):783–794, 2011.
- [4] T. Auckenthaler, H.-J. Bungartz, T. Huckle, L. Krämer, B. Lang, and P. Willems. Developing algorithms and software for the parallel solution of the symmetric eigenvalue problem. *Journal of Computational Science*, 2(3):272–278, 2011.

- [5] U. Baur and P. Benner. Factorized solution of Lyapunov equations based on hierarchical matrix arithmetic. *Computing*, 78(3):211–234, 2006.
- [6] M. Bebendorf. *Hierarchical matrices*, volume 63 of *Lecture Notes in Computational Science and Engineering*. Springer-Verlag, Berlin, 2008. A means to efficiently solve elliptic boundary value problems.
- [7] P. Benner, S. Börm, T. Mach, and K. Reimer. Computing the eigenvalues of symmetric \mathcal{H}^2 -matrices by slicing the spectrum. *Comput. Vis. Sci.*, 16(6):271–282, 2013.
- [8] P. Benner and T. Mach. On the QR decomposition of \mathcal{H} -matrices. *Computing*, 88(3-4):111–129, 2010.
- [9] P. Benner and T. Mach. Computing all or some eigenvalues of symmetric \mathcal{H}_ℓ -matrices. *SIAM J. Sci. Comput.*, 34(1):A485–A496, 2012.
- [10] M. Benzi, P. Boito, and N. Razouk. Decay properties of spectral projectors with applications to electronic structure. *SIAM Rev.*, 55(1):3–64, 2013.
- [11] G. Beylkin, N. Coult, and M. J. Mohlenkamp. Fast spectral projection algorithms for density-matrix computations. *J. Comput. Phys.*, 152(1):32–54, 1999.
- [12] P. Bientinesi, F. D Igual, D. Kressner, M. Petschow, and E. S. Quintana-Ortí. Condensed forms for the symmetric eigenvalue problem on multi-threaded architectures. *Concurrency and Computation: Practice and Experience*, 23(7):694–707, 2011.
- [13] C. H. Bischof, B. Lang, and X. Sun. A framework for symmetric band reduction. *ACM Trans. Math. Software*, 26(4):581–601, 2000.
- [14] Å. Björck. *Numerical Methods for Least Squares Problems*. SIAM, Philadelphia, PA, 1996.
- [15] D. Braess and W. Hackbusch. Approximation of $1/x$ by exponential sums in $[1, \infty)$. *IMA J. Numer. Anal.*, 25(4):685–697, 2005.
- [16] T. A. Davis and Y. Hu. The University of Florida Sparse Matrix Collection. *ACM Transactions on Mathematical Software*, 38(1):1–25, 2011.
- [17] J. W. Demmel, O. A. Marques, B. N. Parlett, and C. Vömel. Performance and accuracy of LAPACK’s symmetric tridiagonal eigensolvers. *SIAM J. Sci. Comput.*, 30(3):1508–1526, 2008.
- [18] I. S. Dhillon, B. N. Parlett, and C. Vömel. The design and implementation of the MRRR algorithm. *ACM Trans. Math. Software*, 32(4):533–560, 2006.
- [19] L. Eldén. Algorithms for the regularization of ill-conditioned least squares problems. *Nordisk Tidskr. Informationsbehandling (BIT)*, 17(2):134–145, 1977.
- [20] L. Eldén. An algorithm for the regularization of ill-conditioned, banded least squares problems. *SIAM J. Sci. Statist. Comput.*, 5(1):237–254, 1984.
- [21] I. P. Gavrilyuk, W. Hackbusch, and B. N. Khoromskij. \mathcal{H} -matrix approximation for the operator exponential with applications. *Numer. Math.*, 92(1):83–111, 2002.
- [22] I. P. Gavrilyuk, W. Hackbusch, and B. N. Khoromskij. Data-sparse approximation to the operator-valued functions of elliptic operator. *Math. Comp.*, 73(247):1297–1324, 2004.

- [23] S. Goedecker. Linear scaling electronic structure methods. *Rev. Mod. Phys.*, 71:1085–1123, 1999.
- [24] G. H. Golub and C. F. Van Loan. *Matrix computations*. Johns Hopkins University Press, Baltimore, MD, fourth edition, 2013.
- [25] L. Grasedyck, W. Hackbusch, and B. N. Khoromskij. Solution of large scale algebraic matrix Riccati equations by use of hierarchical matrices. *Computing*, 70(2):121–165, 2003.
- [26] S. Güttel, E. Polizzi, P. T. P. Tang, and G. Viaud. Zolotarev quadrature rules and load balancing for the FEAST eigensolver. *SIAM J. Sci. Comput.*, 37(4):A2100–A2122, 2015.
- [27] W. Hackbusch. A sparse matrix arithmetic based on \mathcal{H} -matrices. I. Introduction to \mathcal{H} -matrices. *Computing*, 62(2):89–108, 1999.
- [28] W. Hackbusch. *Hierarchical matrices: algorithms and analysis*, volume 49 of *Springer Series in Computational Mathematics*. Springer, Heidelberg, 2015.
- [29] A. Haidar, H. Ltaief, and J. Dongarra. Parallel Reduction to Condensed Forms for Symmetric Eigenvalue Problems Using Aggregated Fine-grained and Memory-aware Kernels. In *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 8:1–8:11. ACM, 2011.
- [30] A. Haidar, H. Ltaief, and J. Dongarra. Toward a high performance tile divide and conquer algorithm for the dense symmetric eigenvalue problem. *SIAM J. Sci. Comput.*, 34(6):C249–C274, 2012.
- [31] A. Haidar, R. Solcà, M. Gates, S. Tomov, T. Schulthess, and J. Dongarra. Leading Edge Hybrid Multi-GPU Algorithms for Generalized Eigenproblems in Electronic Structure Calculations. In *Supercomputing*, volume 7905 of *Lecture Notes in Computer Science*, pages 67–80. Springer Berlin Heidelberg, 2013.
- [32] N. J. Higham and F. Tisseur. A block algorithm for matrix 1-norm estimation, with an application to 1-norm pseudospectra. *SIAM J. Matrix Anal. Appl.*, 21(4):1185–1201 (electronic), 2000.
- [33] M. Lintner. *Lösung der 2D Wellengleichung mittels hierarchischer Matrizen*. Doctoral thesis, TU München, 2002.
- [34] M. Lintner. The eigenvalue problem for the 2D Laplacian in \mathcal{H} -matrix arithmetic and application to the heat and wave equation. *Computing*, 72(3-4):293–323, 2004.
- [35] T. Mach. *Eigenvalue algorithms for symmetric hierarchical matrices*. Doctoral thesis, TU Chemnitz, 2012.
- [36] O. A. Marques, C. Vömel, J. W. Demmel, and B. N. Parlett. Algorithm 880: a testing infrastructure for symmetric tridiagonal eigensolvers. *ACM Trans. Math. Software*, 35(1):Art. 8, 13, 2009.
- [37] Y. Nakatsukasa, Z. Bai, and F. Gygi. Optimizing Halley’s iteration for computing the matrix polar decomposition. *SIAM J. Matrix Anal. Appl.*, 31(5):2700–2720, 2010.
- [38] Y. Nakatsukasa and R. W. Freund. Computing fundamental matrix decompositions accurately via the matrix sign function in two iterations: The power of Zolotarev’s functions. *SIAM Rev.*, 2016. To appear.

- [39] Y. Nakatsukasa and N. J. Higham. Stable and efficient spectral divide and conquer algorithms for the symmetric eigenvalue decomposition and the SVD. *SIAM J. Sci. Comput.*, 35(3):A1325–A1349, 2013.
- [40] P. P. Petrushev and V. A. Popov. *Rational approximation of real functions*, volume 28 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, 1987.
- [41] M. Petschow, E. Peise, and P. Bientinesi. High-performance solvers for dense Hermitian eigenproblems. *SIAM J. Sci. Comput.*, 35(1):C1–C22, 2013.
- [42] H. R. Schwarz. Handbook Series Linear Algebra: Tridiagonalization of a symmetric band matrix. *Numer. Math.*, 12(4):231–241, 1968.
- [43] E. Solomonik, G. Ballard, J. Demmel, and T. Hoefler. A communication-avoiding parallel algorithm for the symmetric eigenvalue problem. arXiv:1604.03703, 2016.
- [44] R. Vandebril, M. Van Barel, and N. Mastronardi. *Matrix computations and semiseparable matrices. Vol. 1*. Johns Hopkins University Press, Baltimore, MD, 2008. Linear systems.
- [45] J. Vogel, J. Xia, S. Cauley, and V. Balakrishnan. Superfast divide-and-conquer method and perturbation analysis for structured eigenvalue solutions. *SIAM J. Sci. Comput.*, 38(3):A1358–A1382, 2016.

Recent publications:

MATHEMATICS INSTITUTE OF COMPUTATIONAL SCIENCE AND ENGINEERING

Section of Mathematics

Ecole Polytechnique Fédérale (EPFL)

CH-1015 Lausanne

- 18.2016** STEFANO PAGANI, ANDREA MANZONI, ALFIO QUARTERONI:
A reduced basis ensemble Kalman filter for state/parameter identification in large-scale nonlinear dynamical systems
- 19.2016** ANDREA MANZONI, FEDERICO NEGRI :
Automatic reduction of PDEs defined on domains with variable shape
- 20.2016** MARCO FEDELE, ELENA FAGGIANO, LUCA DEDÈ, ALFIO QUARTERONI:
A patient-specific aortic valve model based on moving resistive immersed implicit surfaces
- 21.2016** DIANA BONOMI, ANDREA MANZONI, ALFIO QUARTERONI:
A matrix discrete empirical interpolation method for the efficient model reduction of parametrized nonlinear PDEs: application to nonlinear elasticity problems
- 22.2016** JONAS BALLANI, DANIEL KRESSNER, MICHAEL PETERS:
Multilevel tensor approximation of PDEs with random data
- 23.2016** DANIEL KRESSNER, ROBERT LUCE:
Fast computation of the matrix exponential for a Toeplitz matrix
- 24.2016** FRANCESCA BONIZZONI, FABIO NOBILE, ILARIA PERUGIA:
Convergence analysis of Padé approximations for Helmholtz frequency response problems
- 25.2016** MICHELE PISARONI, FABIO NOBILE, PÉNÉLOPE LEYLAND:
A continuation multi level Monte Carlo (C-MLMC) method for uncertainty quantification in compressible aerodynamics
- 26.2016** ASSYR ABDULLE, ONDREJ BUDÁČ:
Multiscale model reduction methods for flow in heterogeneous porous media
- 27.2016** ASSYR ABDULLE, PATRICK HENNING:
Multiscale methods for wave problems in heterogeneous media
- 28.2016** ASSYR ABDULLE, ORANE JECKER:
On heterogeneous coupling of multiscale methods for problems with and without scale separation
- 29.2016** WOLFGANG HACKBUSCH, DANIEL KRESSNER, ANDRÉ USCHMAJEW:
Perturbation of higher-order singular values
- 30.2016** ASSYR ABDULLE, ONDREJ BUDÁČ, ANTOINE IMBODEN:
Multiscale methods and model order reduction for flow problems in three-scale porous media
- 31.2016** DANIEL KRESSNER, ANA ŠUŠNJARA:
Fast computation of spectral projectors of banded matrices