Max-Planck-Institut
für Mathematik
in den Naturwissenschaften
Leipzig

On the Robustness of Elliptic Resolvents Computed by means of the Technique of Hierarchical Matrices

by

Mike Espig, and Wolfgang Hackbusch

Preprint no.: 11 2007

On the Robustness of Elliptic Resolvents Computed by Means of the Technique of Hierarchical Matrices

Mike Espig, Wolfgang Hackbusch
Max-Planck-Institute for Mathematics in the Sciences
Inselstr. 22-26, D-04103 Leipzig, Germany.
espig@mis.mpg.de, wh@mis.mpg.de

Abstract

Hierarchical matrices provide a data-sparse way to approximate fully populated matrices. In this paper we exploit robust $\mathcal{H}$-matrix techniques to approximate the resolvents of stiffness matrices as they appear in (finite element or finite difference) discretisations of elliptic partial differential equations.

Keywords: Hierarchical matrices, data-sparse approximations, resolvents

1 Introduction

Usually, resolvents

$$(z\mathrm{Id} - \mathrm{A})^{-1} \qquad (\mathrm{z} \in \mathbb{C},\ \mathrm{A} \in \mathbb{R}^{\mathrm{n}\times\mathrm{n}})$$

are not often used in numerical applications, since the computation of inverse matrices is prohibitively costly. Even if, for a given right-hand side b, systems of the form $(z\mathrm{Id} - \mathrm{A})\mathrm{x} = \mathrm{b}$ are to be solved for different complex z, the iterative solution method may depend on z. In particular, even if A is positive definite, $A - z\mathrm{Id}$ may be indefinite.

For the interesting case that the matrix A is obtained by a discretisation of an elliptic boundary value problem, the novel technique of hierarchical matrices is able to perform all matrix operations including the inversion approximately with almost linear cost. The exact cost of the inversion is $O(nk^2 \log^2 n)$ (cf. [5]), where k is the local rank (cf. §2).

Due to the almost linear cost, the interest in resolvents is renewed. For instance, the approximation of integrals like

$$\oint_C f(z)(z\mathrm{Id} - \mathrm{A})^{-1}\mathrm{dz} \tag{1.1}$$

over a closed curve C in the complex plane can be used to represent interesting matrix-valued functions. In [4], an open parabola C ($x = cy^2$ for $z = x + iy$) is used to represent the matrix exponential:

$$\exp(-tA) = \frac{1}{2\pi i}\int_C e^{-zt}(z\mathrm{Id} - \mathrm{A})^{-1}\mathrm{dz}. \tag{1.2}$$

In the latter case, the spectrum of A must lie in the interior $\left\{z = x + iy : x > cy^2\right\}$ of the parabola C.

In the case of (1.1), one introduces an appropriate parametrisation. Then the resulting integral with a periodic integrand can be approximated by the trapezoidal rule, which yields exponential convergence when the number N of quadrature points z_i increases. Hence, $N = O(\log^\alpha \frac{1}{\varepsilon})$ evaluations of $(z_i\mathrm{Id} - \mathrm{A})^{-1}$ are sufficient. The parametrisation of the parabola in (1.2) leads to an integral over $(-\infty, +\infty)$, which can be treated by sinc quadrature, yielding a similar size of N (cf. [8, §§2.2-2.4]). We see from these examples that the evaluation of the complex integrals is possible with a cost of $O(n \log^\alpha \frac{1}{\varepsilon} \log^\beta n)$.

The error analysis of the hierarchical matrix operations is done for a fixed $z \in \mathbb{C}$. The dependence of the approximation errors on z and their possibly uniform behaviour is hard to analyse. On the other hand, such results are of fundamental interest for the applications mentioned above. In particular, two questions need to be answered:

- How does the approximation error of the inversion procedure behave when $|z|$ becomes large?
- How does the approximation error behave when z approaches an eigenvalue?

The latter case happens when the curve C intersects the real axis between two real eigenvalues λ_i and λ_{i+1}. Therefore, in [3], the first author investigated the error behaviour with respect to $z \in \mathbb{C}$ *numerically*, at least for A being the discretisation of the Laplace operator. These results are given in Section 3. We first recapitulate some basic facts about hierarchical matrices in Section 2.

2 Hierarchical Matrices

In this section we give a brief introduction to a method for the data-sparse approximation of matrices resulting from the discretisation of non-local operators occurring in boundary integral methods or as the inverses of partial differential operators. The result of the approximation will be so-called hierarchical matrices (or short $\mathcal{H}$-matrices, c.f. [2], [5], [7]). These matrices form a subset of the set of all matrices and have a data-sparse representation. The essential operations for these matrices (matrix-vector and matrix-matrix multiplication, addition and inversion) can be approximately performed in linear logarithmic complexity, [5].
In §3 the index set I consists of the finite element nodal points. The formal definition of an $\mathcal{H}$-matrix depends on appropriate hierarchical partitionings of the index set and also of the product index set which are organised in (block) cluster trees as defined below.

Definition 2.1 (cluster tree) *Let I be a finite set, and let $T(I) = (V, E)$ be a tree with vertex set V and edge set E. For a vertex $v \in V$ we define the set of sons of v as $S(v) := \{w \in V : (v, w) \in E\}$. The tree $T(I)$ is called a cluster tree of I if its vertices consist of subsets of I and satisfy the following conditions:*

- *$I \in V$ is the root of $T(I)$ and $v \subset I$, $v \neq \emptyset$, for all $v \in V$.*
- *For all $v \in V$ there either holds $S(v) = \emptyset$ or $v = \dot{\cup}_{w \in S(v)} w$.*

In the following we identify V and $T(I)$; i.e., we write $v \in T(I)$ instead of $v \in V$. The nodes $v \in V$ are called clusters. The support of a cluster $v \in T(I)$ is given by the union of the supports of the basis functions corresponding to its elements, i.e.,

$$\Omega_v := \bigcup_{i \in v} \Omega_i, \quad \text{where } \Omega_i := \text{supp } \varphi_i \text{ for all } i \in I.$$

Definition 2.2 (leaf, father, level, depth) *Let $T(I)$ be a cluster tree. The set of leaves of the tree $T(I)$ is $\mathcal{L} := \{v \in T(I) : S(v) = \emptyset\}$. The uniquely determined predecessor (father) of a nonroot vertex $v \in T(I)$ is denoted by $\mathcal{F}(v)$. The levels of the tree $T(I)$ are defined by*

$$T_I^{(0)} := \{I\}, \quad T_I^{(l)} := \{v \in T(I) : \mathcal{F}(v) \in T_I^{(l-1)}\}, \quad \text{for } l \in \mathbb{N},$$

and we write level$(v) = l$ if $v \in T_I^{(l)}$. The depth of $T(I)$ is defined as $d(T) := \max\{l \in \mathbb{N} : T_I^{(l)} \neq \emptyset\}$.

Definition 2.3 (block cluster tree) *Let $T(I)$ be a cluster tree of the index set I. A cluster tree $T(I \times I)$ is called a block cluster tree (based upon $T(I)$) if for all $v \in T_{I \times I}^{(l)}$ there exist $t, s \in T_I^{(l)}$ such that $v = t \times s$. The nodes $v \in T(I \times I)$ are called block clusters.*

Analogously to the cluster tree, for any block cluster tree $T(I \times I)$ there holds $I \times I = \dot{\bigcup}\{v : v \in \mathcal{L}(T_{I\times I})\}$; i.e., the leaves of the block cluster tree provide a disjoint block partition of the product index set $I \times I$. The objective is to construct a block cluster tree from a given cluster tree such that the leaves (of the block cluster tree) correspond to (preferably large) matrix blocks with smooth data that can be approximated by low rank matrices, [2]. Figure 2.1 shows a partition for a 1D case where I corresponds to nodal points in an interval. [1]

In the following construction, we will build a block cluster tree iteratively starting from $I \times I$ and

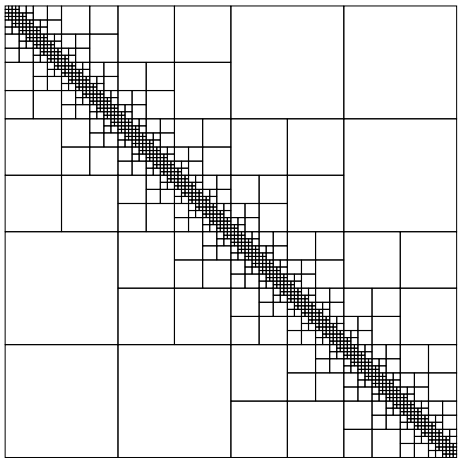

Figure 2.1: Example of a matrix partition

refining the block clusters if they do not satisfy a certain admissibility condition. The choice of the admissibility condition depends on the underlying continuous problem (i.e., elliptic partial differential equation, in particular its associated Greens function) and shall ensure that all admissible blocks allow a sufficiently accurate low rank approximation. A typical admissibility condition for uniformly elliptic problems is as follows:

$$\text{block cluster } s \times t \text{ is admissible } :\Leftrightarrow \min(\operatorname{diam}(\Omega_s), \operatorname{diam}(\Omega_t)) \leq \eta \operatorname{dist}(\Omega_s, \Omega_t),$$

where Ω_s and Ω_t are the supports of the clusters, see definition 2.1.
A given cluster tree along with an admissibility condition allows the following canonical construction of a block cluster tree, [5]:
Let the cluster tree $T(I)$ be given. We define the block cluster tree $T(I \times I)$ by $\operatorname{root}(T) := I \times I$, and each vertex $s \times t \in T(I \times I))$ has the set of successors:

$$S(t \times s) := \begin{cases} \emptyset & \text{if } s \times t \text{ is admissible,} \\ \emptyset & \text{if } \min\{\#s, \#t\} \leq n_{\min}, \\ \{\tilde{s} \times \tilde{t} : \tilde{s} \in S(s), \tilde{t} \in S(t)\} & \text{otherwise .} \end{cases}$$

The parameter $n_{\min}$ ensures that blocks do not become so small that the matrix arithmetic of a full matrix is more efficient than any further subdivision. It is typically set to $n_{\min} = 16$ or even $n_{\min} = 32$.

Definition 2.4 ($\mathcal{H}$-matrix) *Let $k, n_{\min} \in \mathbb{N}_0$. The set of $\mathcal{H}$-matrices induced by a block cluster tree $T := T(I \times I)$ with blockwise rank k and minimum block size $n_{\min}$ is defined by*

$$\mathcal{H}(T, k) := \{M \in \mathbb{C}^{I\times I} : \forall t \times s \in \mathcal{L}(T) : rank(M|_{t\times s}) \text{ or } \min\{\#s, \#t\} \leq n_{\min}\}.$$

[1] For an illustration, a 1D example is more appropriate since any illustration must order the indices in a subsequent order. However, only in the 1D case there is a natural ordering.

For the practical use it is essential that each matrix block $A|_b$ with $b = t \times s$ is represented as a product $B_b * C_b^\top$ with $B_b \in \mathbb{C}^{t\times\{1,\dots,k\}}$ and $C_b \in \mathbb{C}^{s\times\{1,\dots,k\}}$. In practice, matrix blocks for $b = t \times s$ of sufficiently small size, $\min\{\#t, \#s\} \le n_{\min}$ are stored instead as full matrices, as described above.

The first approximate operation is the matrix addition, where $A, B \in \mathcal{H}(T, k)$ leads to $rank\left((A + B)|_b\right) \le 2k$. For instance by SVD the exact sum $(A + B)|_b$ is replaced by the optimal approximation of rank k, producing an approximation in $\mathcal{H}(T, k)$. Due to the hierarchical structure of block-cluster tree the multiplication $A * B$ can be performed recursively including truncations of the matrix blocks back to rank k. Inversion by means of the block Gauss elimination requires only the foregoing matrix operations $\pm, *$ and recursively accomplished inversions. For a more detail description see also [2], [5] and [7].

Since the diagonal entries belong to the small blocks (size $\le n_{\min}$) which are treated as full matrices (see Figure 2.1), a simple observation can be stated.

Remark 2.5 *If $A \in \mathcal{H}(T, k)$, then $z\mathrm{Id} - \mathrm{A}$ also belongs to $\mathcal{H}(T, k)$, for all $z \in \mathbb{C}$.*

Due to the choice of the admissibility condition, the following statement is a simple consequence.

Remark 2.6 *Any finite-element matrix A belongs to $\mathcal{H}(T, k)$ for all $k \in \mathbb{N}_0$, since $A|_b = 0$ for all $b = t \times s \in T$ with $\min\{\#t, \#s\} > n_{\min}$.*

A more difficult question is whether the (exact) inverse A^{-1} of a finite-element matrix A has good approximants in $\mathcal{H}(T, k)$. The answer is positive, and, astonishingly, the result does not depend on the smoothness of the coefficients in the differential operator (details in [1]).

Theorem 2.7 ([1]) *Let $Lu = f$ be a boundary value problem in a Lipschitz domain $\Omega \subset \mathbb{R}^d$ with $L = \operatorname{div} K(x) \operatorname{grad}$, where the symmetric $d \times d$-matrix $K(\cdot) \in L^\infty(\Omega)$ has eigenvalues in the interval $[a, b]$ with $0 < a \le b < \infty$. Then the inverse A^{-1} of a finite-element matrix has approximants in $\mathcal{H}(T, k)$ so that the error decays exponentially in k.*

Since $z\mathrm{Id} - \mathrm{A}$ corresponds to a compact perturbation of $-A$, the same result holds for the resolvent $(z\mathrm{Id} - \mathrm{A})^{-1}$ for any regular value of z. In further discussions we will denote the approximated inverse of A in $\mathcal{H}(T, k)$ by $inv_{\mathcal{H}}(A)$.

3 Numerical Results

3.1 Model Problem

Theorem 2.7 indicates that hierarchical matrix inversion is robust with respect to the PDE coefficients and to the domain. This is confirmed by numerical experiments. Hence, it suffices to test the resolvents for the matrix A obtained from the Poisson equation in a square.

Definition 3.1 (Model Problem) *The* model problem *is the Poisson equation $-\triangle u = f$ with Dirichlet condition $u|_{\partial\Omega} = 0$ in the square $\Omega := (0, 1)^2 \subset \mathbb{R}^2$. The finite element discretisation uses the uniform grid from Figure 3.1 and piecewise linear elements. The (positive definite) finite element matrix is denoted by A. The step size is $h = 1/N$ so that the size of the matrix is*

$$n = (N - 1)^2. \tag{3.1}$$

For the following numerical experiments the knowledge of all eigenvalues of the matrix A is helpful.

Remark 3.2 *The matrix A from Definition 3.1 has the eigenvalues*

$$\lambda_{\ell m} = 4\left[\sin^2\left(\frac{\ell\pi h}{2}\right) + \sin^2\left(\frac{m\pi h}{2}\right)\right] \qquad (1 \le \ell, m \le N - 1),$$

which are not all different. The multiplicity of $\lambda_{\ell m}$ is given by the number of equal indices $\lambda_{\ell m} = \lambda_{\ell' m'}$ $(1 \le \ell', m' \le N - 1)$. The extreme eigenvalues are

$$\lambda_{\min} = \lambda_{1,1} = 8\sin^2\left(\frac{\pi h}{2}\right) = \|A^{-1}\|_2^{-1}, \qquad \lambda_{\max} = \lambda_{N-1,N-1} = 8\cos^2\left(\frac{\pi h}{2}\right) = \|A\|_2.$$

Proof. See [6, Thm. 4.1.1 and Lemma 4.1.2]. ∎

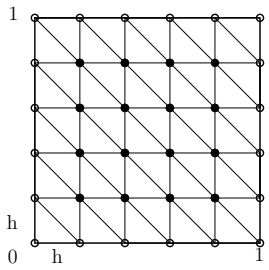

Figure 3.1: Regular finite element triangulation of Ω

3.2 Error Consideration on Selected Curves

We have calculated absolute errors for the resolvents in Frobenius norm $\|\cdot\|_F$.

Notation 3.3 (Approximation Error) *Let $A \in \mathbb{R}^{n\times n}$ be the matrix from Definition 3.1. The error of the resolvent is denoted by*

$$\triangle_F^n(z) := \|Id - inv_{\mathcal{H}}(z\mathrm{Id} - \mathrm{A})(\mathrm{zId} - \mathrm{A})\|_{\mathrm{F}} \qquad \textit{for regular values } \mathrm{z} \in \mathbb{C}.$$

We use $\triangle_F^n$ instead of $\triangle_F^n(z)$, if the z-dependence is clear from the context.

The following plots represent the approximation error in logarithmic scaling. The ordinates range from minimal to maximal error value. On the abscissae the parameter interval of the curve is visible.

3.2.1 Curve between two Eigenvalues

First, we consider a curve which intersects the real axis and therefore comes close to two neighbouring eigenvalues. In this case, the curve meets the real axis between $\lambda_{10,10} = 0.679772$ and $\lambda_{09,11} = 0.690279$. Its definition is

$$\gamma_1(\tau) := i \cdot \tau + \frac{1}{2}(\lambda_{10,10} + \lambda_{09,11}) \qquad \text{for } 0 \le \tau \le 1$$

(see also Figure 3.2).

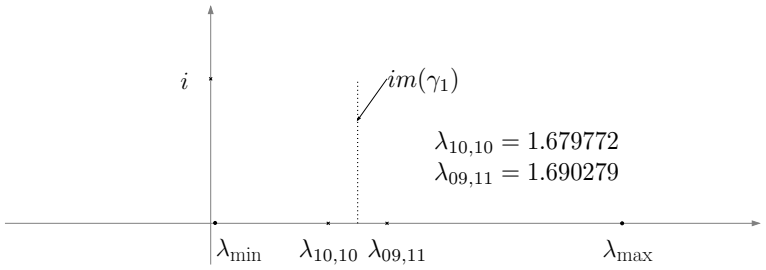

Figure 3.2: Image of γ_1, for $0 \le \tau \le 1$.

The discretisation parameter is $n = 1024$. The minimum and maximum value of the error

$\triangle_F^{1024}(\gamma_1(\tau))$, $0 \le \tau \le 1$, is shown in the following Table for different choices of the local rank k:

local rank k	$\min(\triangle_F^{1024})$	$\max(\triangle_F^{1024})$
4	2.9×10^{-5}	9.8×10^{-4}
5	7.0×10^{-6}	9.5×10^{-5}
6	3.0×10^{-6}	1.4×10^{-5}
7	2.0×10^{-6}	4.0×10^{-6}
8	1.0×10^{-6}	2.0×10^{-6}
9	6.5×10^{-8}	0.9×10^{-7}
10	4.0×10^{-10}	8.5×10^{-10}
11	1.5×10^{-14}	6.0×10^{-14}

Obviously, the minimum and maximum values are not so different. The (exponential) decrease of the error with respect to k is the typical behaviour. For $k = 9$ and $k = 10$ the error dependence of the curve parameter τ is shown in Figure 3.3.

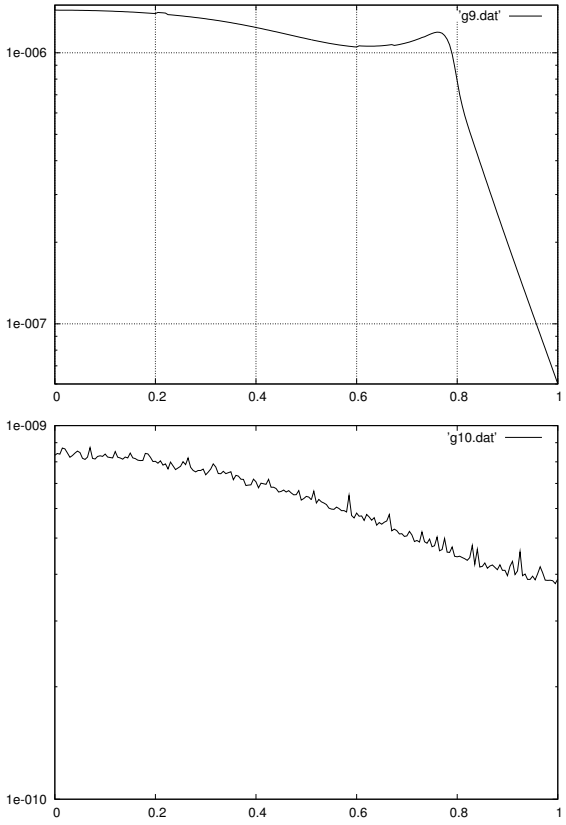

Figure 3.3: Behaviour of the approximation error $\triangle_F^{1024}$ as a function of τ on $\gamma_1(\tau) := i \cdot \tau + \frac{1}{2}(\lambda_{10,10} + \lambda_{09,11})$. *Left:* $k = 9$, *right:* $k = 10$.

3.2.2 Curve intersecting an Eigenvalue

The previous example corresponds to the desired case that the curve passing the gap between two eigenvalues should have similar distances from both eigenvalues. If the eigenvalues are not known, or only approximately known, it may happen that the curve comes very close to one of the eigenvalues. The same happens if there is a dense cluster of eigenvalues. The difficulties of inverting $z\mathrm{Id} - \mathrm{A}$ frequently occur in eigenvalue problems with shift, where z may be very close to an eigenvalue of A. Our method can be also used in these cases.

Therefore, next we consider the extreme case of a curve intersecting the real axis at the eigenvalue $\lambda_{3,3} = 0.1620281\ldots$ The definition of the curve γ_2 in this test is

$$\gamma_2(\tau) = i \cdot \tau + \lambda_{3,3} \qquad \text{for } 0 < \tau \leq 1.$$

By definition, the resolvent becomes singular at $\tau = 0$. Therefore, it is not astonishing that the error $\triangle_F^{1024}(\gamma_2(\tau))$ increases as τ approaches 0. But, as can be seen from Figure 3.4 for $k = 9$ and Figure 3.5 for $k = 10$, the error has a stable behaviour until $\tau \approx 5 \times 10^{-3}$. Only in a tiny neighbourhood of the eigenvalue does the computation of a proper resolvent approximation fail.

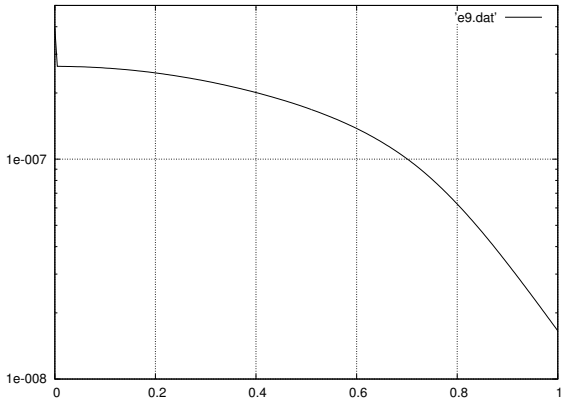

Figure 3.4: Behaviour of the approximation error $\triangle_F^{1024}$ as a function of τ on $\gamma_2(\tau) := i \cdot \tau + \lambda_{3,3}$ for $k = 9$.

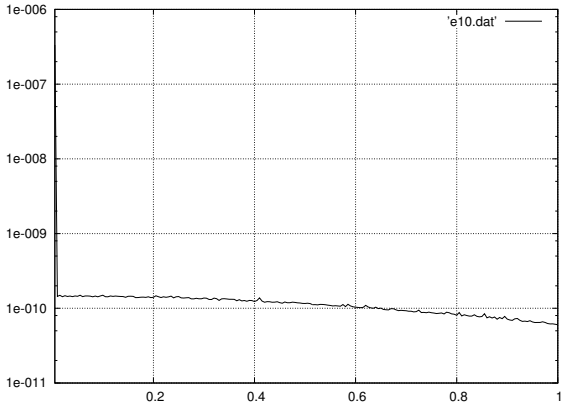

Figure 3.5: Behaviour of the approximation error $\triangle_F^{1024}$ as a function of τ on $\gamma_2(\tau) := i \cdot \tau + \lambda_{3,3}$ for $k = 10$.

At the first glance, it may be surprising that the computation of the resolvents can be performed in a robust way close to the singularity. The reason is stated in the following remark.

Remark 3.4 *Assume that $A = A^H$. Let z_0 be a simple eigenvalue corresponding to the normalised eigenvector v. When z approaches z_0, the resolvent $(z\mathrm{Id} - \mathrm{A})^{-1}$ behaves like $R_0(z) + \frac{1}{z-z_0} vv^H$, where $R_0(z)$ is continuous (even analytic) at z_0. Due to [1], $R_0(z)$ can be well approximated by hierarchical matrices. The term $\frac{1}{z-z_0} vv^H$ is large, but its restriction to each block $b \in T$ (see §2) is a rank-1 matrix which can be represented exactly.*

Generalisations to non-Hermitean matrices or multiple eigenvalues are straightforward. For non-Hermitean matrices we can consider the singular value decomposition of A ($R_0(z) + \frac{1}{z-z_0} vv^H$ is

replaced by $R_0(z) + \frac{1}{z-z_0} uv^H$). In the multiple eigenvalue case the restriction to each block can be represented by a rank-k matrix, where k is the multiplicity of the eigenvalue.

3.2.3 Curves with $|z| \to \infty$

The example (1.2) is an integral over a parabola, which can be considered as a closed curve intersecting $z = \infty \in \bar{\mathbb{C}}$. Therefore, we choose the curve

$$\gamma_P(\tau) := \tau^2 + i\tau.$$

For the local ranks $k = 4, 10$ the errors $\triangle_F^{1024}$ are plotted in Fig. 3.6. The maximum error $\max(\triangle_F^{1024})$ is attained in the interval $[-2, 2]$. For $\tau > 7$ ($k = 4$), $\tau > 3$ ($k = 10$) and in particular for $\tau \to \pm\infty$ the error is smaller than 10^{-14}. The approximation error decreases for $|\tau| > 2$ very fast in both cases.

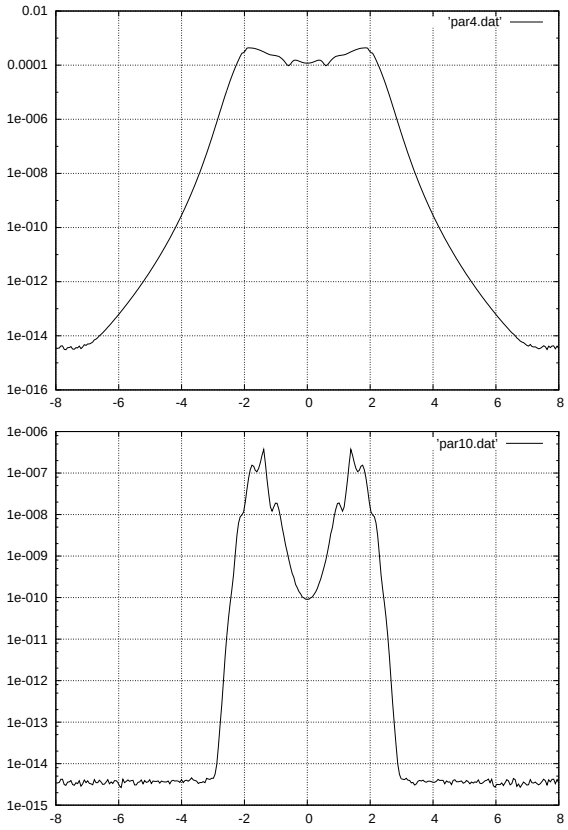

Figure 3.6: Behaviour of the approximation error $\triangle_F^{1024}$ as a function of τ on $\gamma_p(\tau) := \tau^2 + i\tau$. *Left:* $k = 4$, *right:* $k = 10$.

4 Conclusion

The computation of the resolvents appears to be rather stable. Large values of z are not dangerous as $(z\mathrm{Id} - \mathrm{A})^{-1}$ becomes small anyway. More remarkably, the computation of $(z\mathrm{Id} - \mathrm{A})^{-1}$ in the neighbourhood of singular values z_0 only fails if $|z - z_0|$ is very small (5×10^{-3} in the example of §3.2.2). As a consequence, given a curve γ for a matrix with unknown spectrum, it is rather unlikely that γ comes so close to one of the eigenvalues. The more severe drawback is that the convergence of the quadrature error slows down if there are nearby singularities, when we replace the exact integral by a quadrature formula.

References

[1] M. Bebendorf and W. Hackbusch. Existence of $\mathcal{H}$-matrix approximants to the inverse FE-matrix of elliptic operators with L^∞-coefficients. Numer. Math., 95:1–28 (2003).

[2] S. Börm, L. Grasedyck, and W. Hackbusch. Introduction to hierarchical matrices with applications. Eng. Anal. Bound. Elem., 27:405–422 (2003).

[3] M. Espig. Approximation von Resolventen des Laplace-Operators mit Hilfe hierarchischer Matrizen. Diplom thesis, Christian-Albrechts-Universität zu Kiel (2004).

[4] I. P. Gavrilyuk, W. Hackbusch, and B. Khoromskij. $\mathcal{H}$-matrix approximation for the operator exponential with applications. Numer. Math., 92:83–111 (2002).

[5] L. Grasedyck and W. Hackbusch. Construction and arithmetics of $\mathcal{H}$-matrices. Computing, 70:295–334 (2003).

[6] W. Hackbusch. Iterative solution of large sparse systems of equations. Springer-Verlag, New York (1994).

[7] W. Hackbusch. A sparse matrix arithmetic based on $\mathcal{H}$-matrices. Part I: Introduction to $\mathcal{H}$-matrices. Computing, 62:89–108 (1999).

[8] W. Hackbusch and B. Khoromskij. Low-rank Kronecker-product approximation to multi-dimensional nonlocal operators. Part I. Separable approximation of multi-variate functions. Computing, 76:177–202 (2006).