

A nested Krylov subspace method to compute the sign function of large complex matrices[☆]

Jacques C.R. Bloch, Simon Heybrock

Institute for Theoretical Physics, University of Regensburg, 93040 Regensburg, Germany

Abstract

We present an acceleration of the well-established Krylov-Ritz methods to compute the sign function of large complex matrices, as needed in lattice QCD simulations involving the overlap Dirac operator at both zero and nonzero baryon density. Krylov-Ritz methods approximate the sign function using a projection on a Krylov subspace. To achieve a high accuracy this subspace must be taken quite large, which makes the method too costly. The new idea is to make a further projection on an even smaller, nested Krylov subspace. If additionally an intermediate preconditioning step is applied, this projection can be performed without affecting the accuracy of the approximation, and a substantial gain in efficiency is achieved for both Hermitian and non-Hermitian matrices. The numerical efficiency of the method is demonstrated on lattice configurations of sizes ranging from 4^4 to 10^4 , and the new results are compared with those obtained with rational approximation methods.

1. Introduction

In quantum chromodynamics (QCD) some physical observables rely on the chiral properties of the theory. To study such observables in a lattice formulation of QCD it is important to discretize the Dirac operator such that it respects the corresponding chiral symmetry. This is most faithfully achieved using the overlap Dirac operator [1, 2]. To study QCD at nonzero baryon density the overlap formulation was recently extended to include a quark chemical potential [3, 4]. A major ingredient in the overlap operator, which makes its use very challenging, is the computation of the sign function of a complex matrix, which is Hermitian at zero baryon density, but becomes non-Hermitian when a quark chemical potential is introduced.

The search for efficient numerical methods to compute the sign function for the large sparse matrices encountered in this context is an ongoing field of research. Typically, Krylov subspace methods are employed to evaluate the operation of a matrix function on an arbitrary vector. We distinguish two main variants: the Krylov-Ritz approximation, which evaluates the function via a projection on the Krylov subspace, and the rational approximation, where the function is first approximated by a partial fraction expansion, which is then efficiently solved using a multi-shift Krylov subspace inverter.

In the Hermitian case efficient rational approximation methods for the sign function have been devised [5, 6] and are currently being used in large scale lattice simulations. The current method of choice uses the Zolotarev partial fraction expansion [6, 7, 8], which yields the optimal rational approximation to the sign function over a real interval [9], in conjunction with a multi-shift conjugate gradient inversion. For non-Hermitian matrices, which occur in the presence of a quark chemical potential, Krylov subspace approximations to the sign function are relatively new and still under development. Recently, partial fraction expansion methods using the Neuberger expansion [5] with non-Hermitian multi-shift inverters were proposed [10].

The Krylov-Ritz approximation, which we discuss in this paper, is based on the construction of a Krylov basis and its accompanying Ritz matrix. Depending on the algorithm used to construct the basis we distinguish between the Lanczos approximation in the Hermitian case [6], and the Arnoldi approximation [11] or two-sided Lanczos

[☆]Supported by the DFG collaborative research center SFB/TR-55 “Hadron Physics from Lattice QCD”.

approximation [12] in the non-Hermitian case. The latter clearly yields the more efficient function approximation for non-Hermitian matrices [12]. In the Krylov-Ritz approximation the large complex matrix is projected on the Krylov subspace, and its sign function is approximated by lifting the sign function of its projected image (Ritz matrix) back to the original space. The latter sign function is computed to high accuracy using the spectral definition of a matrix function or using a matrix-iterative method. When a large Krylov subspace is needed to reach the desired accuracy, the computation of this matrix sign function becomes a bottleneck for the algorithm.

Herein we will introduce an enhancement of the Krylov-Ritz approximation method which substantially reduces the cost of this internal sign computation and boosts the efficiency of the overall method, such that it competes with, and even surpasses, the rational function approximation in both the Hermitian and non-Hermitian case. The dramatic reduction in computation time is achieved by projecting the Ritz matrix on an even smaller, nested Krylov subspace, after performing a suitable preconditioning step first. The desired sign function is then computed via the sign function of the inner Ritz matrix, which yields the same accuracy as the original Krylov-Ritz approximation.

The outline of the paper is as follows. In Sec. 2 we introduce the overlap operator and the matrix sign function. In Sec. 3 we show how the matrix function of large matrices is computed using Krylov-Ritz approximation methods. In Sec. 4 we introduce the nested Krylov subspace method, which substantially enhances the efficiency of the Krylov-Ritz approximation to the sign function. We study its convergence properties and present numerical results for various lattice sizes, including a comparison with rational approximation methods. Finally, our conclusions are given in Sec. 5. For completeness we have added some algorithms in Appendix.

2. Overlap operator and the matrix sign function

Our motivation to develop numerical algorithms to compute the matrix sign function of large, sparse, complex matrices comes from its application in lattice quantum chromodynamics (LQCD). The overlap formulation of the Dirac operator [1, 2], which ensures that chiral symmetry is preserved in LQCD, is given in terms of the matrix sign function [13], and its definition in the presence of a quark chemical potential μ [3] is given by

$$D_{\text{ov}}(\mu) = \mathbb{1} + \gamma_5 \text{sgn}(\gamma_5 D_w(\mu)), \quad (1)$$

where $\mathbb{1}$ denotes the identity matrix, $\gamma_5 = \gamma_1 \gamma_2 \gamma_3 \gamma_4$ with $\gamma_1, \dots, \gamma_4$ the Dirac gamma matrices in Euclidean space, sgn is the matrix sign function, and

$$D_w(\mu) = \mathbb{1} - \kappa \sum_{i=1}^3 (T_i^+ + T_i^-) - \kappa (e^{\mu} T_4^+ + e^{-\mu} T_4^-) \quad (2)$$

is the Wilson Dirac operator at nonzero chemical potential [14] with $(T_v^\pm)_{yx} = (\mathbb{1} \pm \gamma_v) U_{x,\pm v} \delta_{y,x \pm \hat{v}}$, $\kappa = 1/(8 + 2m_w)$, $m_w \in (-2, 0)$ and $U_{x,\pm v} \in \text{SU}(3)$, where $U_{x,-v} = U_{x-\hat{v},v}^\dagger$. The exponential factors $e^{\pm\mu}$ implement the quark chemical potential on the lattice. For $\mu = 0$ the argument of the sign function is Hermitian, while for $\mu \neq 0$ it is non-Hermitian. To compute the overlap operator we need to define the matrix sign function for a general complex matrix A of dimension n . A generic matrix function $f(A)$ can be defined by

$$f(A) = \frac{1}{2\pi i} \oint_{\Gamma} f(z) (zI - A)^{-1} dz, \quad (3)$$

where Γ is a collection of contours in \mathbb{C} such that f is analytic inside and on Γ and such that Γ encloses the spectrum of A . If A is diagonalizable, i.e., $A = U\Lambda U^{-1}$, with diagonal eigenvalue matrix $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$ and $U \in \text{GL}(n, \mathbb{C})$, then this general definition can be simplified to the well-known spectral form

$$f(A) = U f(\Lambda) U^{-1}, \quad (4)$$

with

$$f(\Lambda) = \text{diag}(f(\lambda_1), \dots, f(\lambda_n)). \quad (5)$$

If A cannot be diagonalized, a spectral definition of $f(A)$ can still be derived using the Jordan decomposition [15]. For simplicity, but without loss of generality, we assume diagonalizability in the following. For Hermitian A the eigenvalues are real and their sign is defined by $\text{sgn}(x) = \pm 1$ for $x \gtrless 0$ with $x \in \mathbb{R}$, such that Eq. (4) readily defines the matrix sign function. For non-Hermitian A the eigenvalues are complex and require a definition of $\text{sgn}(z)$ for $z \in \mathbb{C}$. The sign function needs to satisfy $(\text{sgn}(z))^2 = 1$ and reproduce the usual $\text{sgn}(x)$ for real x . We define

$$\text{sgn}(z) = \frac{z}{\sqrt{z^2}} = \text{sgn}(\text{Re}(z)), \quad (6)$$

where the cut of the square root is chosen along the negative real axis. This choice, although not unique, gives the correct physical result for the overlap Dirac operator in Eq. (1) (see Ref. [4]).

3. Krylov-Ritz approximations for matrix functions

Since we aim at problems with large matrices, as is the case in LQCD, memory and computing power limitations require sophisticated methods to deal with the sign function. For a matrix A of large dimension n the common approach is not to compute $f(A)$ but rather its action on a vector, i.e., $y = f(A)x$, which is needed by iterative inverters to compute $f(A)^{-1}b$ or by iterative eigenvalues solvers for $f(A)$. The Krylov-Ritz method approximates the resulting vector in the Krylov subspace

$$\mathcal{K}_k(A, x) \equiv \text{span}(x, Ax, A^2x, \dots, A^{k-1}x) \quad (7)$$

of \mathbb{C}^n , implicitly making a polynomial approximation of degree $k-1$ to $f(A)$. The optimal approximation to y in this subspace is its orthogonal projection y_k^\perp . For $V_k = (v_1, \dots, v_k)$, where the v_i form an orthonormal basis of $\mathcal{K}_k(A, x)$, an orthogonal projector is given by $P = V_k V_k^\dagger$, and we have

$$y = f(A)x \approx y_k^\perp = P f(A)x. \quad (8)$$

However, to compute this projection on the Krylov subspace we already need y , which is the quantity we wanted to determine in the first place. Thus, we need to replace this exact projection by an approximation. To reduce the large dimensionality of the problem one typically projects A on the Krylov subspace using $A_k \equiv PAP$. The projected matrix A_k has dimension n but rank at most k . The k -dimensional image of the projected matrix A_k is defined by the matrix $H_k = V_k^\dagger A V_k$, which is often referred to as Ritz matrix. The components of H_k are the projection coefficients of A_k in the basis V_k , as A_k and H_k are related by $A_k = V_k H_k V_k^\dagger$ (in analogy to the vector case).

The Krylov-Ritz approximation [16, 17] to $f(A)$ consists in taking the function of the Ritz matrix H_k and lifting it back to the full n -dimensional space,

$$f(A) \approx V_k f(H_k) V_k^\dagger. \quad (9)$$

This approximation actually replaces the polynomial interpolating f at the eigenvalues of A by the polynomial interpolating f at the eigenvalues of H_k , also called Ritz values [17]. Substituting the approximation (9) in $f(A)x$ yields

$$y \approx V_k f(H_k) V_k^\dagger x = |x| V_k f(H_k) e_1^{(k)}, \quad (10)$$

where we choose v_1 collinear with x , i.e., $v_1 = V_k e_1^{(k)} \equiv x/|x|$, with $e_1^{(k)}$ the first unit vector of \mathbb{C}^k . To evaluate the approximation (10) we do not need to perform the matrix multiplications of Eq. (9) explicitly. First, one computes the function $f(H_k)$ of the k -dimensional Ritz matrix to high accuracy, using the spectral definition (4) or a matrix-iterative method. Then, the final approximation is simply a linear combination of the basis vectors v_i , with coefficients given by the first column of $f(H_k)$ multiplied with $|x|$.

The Krylov-Ritz approximation described above uses an orthonormal basis of $\mathcal{K}_k(A, x)$. For the Hermitian case such a basis can efficiently be constructed using the Lanczos algorithm, which we listed in Appendix A for completeness. It generates an orthonormal basis and a tridiagonal symmetric H_k using a three-term recurrence relation. The non-Hermitian case is more laborious as the construction of an orthonormal basis is typically performed using the Arnoldi algorithm, which suffers from long recurrences as each basis vector has to be orthogonalized with respect

to all the previous ones. The two-sided Lanczos algorithm is a suitable alternative [12] which uses two three-term recurrences to construct bases $V_k = (v_1, \dots, v_k)$ and $W_k = (w_1, \dots, w_k)$ of the right, respectively left, Krylov subspaces $\mathcal{K}_k(A, x)$ and $\mathcal{K}_k(A^\dagger, x)$, which are biorthonormal, i.e., $v_i^\dagger w_j = \delta_{ij}$ (see Appendix B for a listing of the algorithm). The lack of orthogonality of the basis V_k prevents the construction of the orthogonal projector needed for the Krylov-Ritz function approximation (9). Nevertheless, the biorthonormality between V_k and W_k can be used to construct an oblique projector $P = V_k W_k^\dagger$ on the right Krylov subspace. The oblique projection of A is $A_k = PAP$ and its k -dimensional image is defined by $H_k = W_k^\dagger A V_k$, which we call two-sided Ritz matrix, such that $A_k = V_k H_k W_k^\dagger$. The matrix H_k generated by the two-sided Lanczos algorithm is tridiagonal. The two-sided Krylov-Ritz approximation to $f(A)$ then consists in taking the matrix function of H_k and lifting it back to the original space,

$$f(A) \approx V_k f(H_k) W_k^\dagger. \quad (11)$$

After applying this approximation of $f(A)$ to x we find an expression which is similar to Eq. (10),

$$y \approx V_k f(H_k) W_k^\dagger x = |x| V_k f(H_k) e_1^{(k)}, \quad (12)$$

where the last step assumes that $v_1 = V_k e_1^{(k)} \equiv x/|x|$. The price paid to achieve short recurrences in the non-Hermitian case is the loss of orthogonality of the projection on the Krylov subspace, which translates in a somewhat lower accuracy of the two-sided Lanczos approximation compared to the Arnoldi approximation, for equal Krylov subspace sizes. Nevertheless, the large gain in speed makes it by far the more efficient method [12].

In the case where f is the sign function, the approximations (10) and (12) require the computation of $\text{sgn}(H_k)$. Although it could be computed directly with the spectral definition (4), matrix-iterative methods are often cheaper for medium sized matrices. We choose to employ the Roberts-Higham iteration (RHi) [18]: Set $S_0 = H_k$ and compute

$$S_{n+1} = \frac{1}{2}(S_n + S_n^{-1}). \quad (13)$$

This iteration converges quadratically to $\text{sgn}(H_k)$, if the sign function for complex arguments is defined by Eq. (6). The matrix inversion scales like k^3 and so will the RHi. For the QCD application considered here, typically 7 to 10 iterations are necessary to converge within machine precision [11, 12].

The hope is that the Krylov-Ritz approximations (10) and (12) are accurate for $k \ll n$. The method is known to work very well as long as no eigenvalues are close to a function discontinuity. However, for the sign function this method suffers from the sign discontinuity along the imaginary axis. If A has eigenvalues close to this discontinuity the approximating polynomial must steeply change from -1 to $+1$ over a small interval to give an accurate approximation. This cannot be achieved with a low order polynomial, i.e., the Krylov subspace must be large, which makes the algorithm expensive. The common solution to this problem is to use deflation, where the contribution of the eigencomponents associated to these critical eigenvalues to the sign function is computed exactly.¹ The Krylov subspace approximation is then performed in a deflated space, i.e., the subspace where the directions along the critical eigenvectors have been removed. We refer to the literature for details [11].

The convergence of the Krylov-Ritz approximations to the matrix sign function is illustrated in Fig. 1: the Lanczos approximation for the Hermitian case on the left, and the two-sided Lanczos approximation for the non-Hermitian case on the right. The accuracy of the approximation cannot be determined by comparing to the exact value $\text{sgn}(A)x$, as its evaluation by direct methods is too costly if A is large. To obtain an estimate for the error, we compute $\tilde{x} \approx \text{sgn}(A)^2 x$ (by applying the Krylov-Ritz approximation twice in succession), which should equal x if the approximation to the sign function were exact, and then take $\varepsilon = |\tilde{x} - x|/2|x|$ as a measure for the error. This error estimate proved to be consistent with the true error obtained by comparing the approximation to the exact solution for 4^4 and 6^4 lattices, and will therefore be used for all lattice sizes. Here, and in all subsequent tests, we choose the test vector $x = (1, \dots, 1)$. As expected, the accuracy improves with increasing Krylov subspace size k , and a larger deflation gap Δ , corresponding to a higher number of deflated eigenvectors, leads to a faster convergence. For a given accuracy and equal deflation gap, the subspace size k required for non-Hermitian A is larger than for Hermitian A .

¹In practice we deflated the eigenvalues with smallest modulus $|\lambda|$ instead of those with smallest absolute real part $|\text{Re } \lambda|$, as the former are more efficiently determined numerically, and both choices yield almost identical deflations for the operator $\gamma_5 D_w(\mu)$ of Eq. (1). The reason for this is that, as long as the chemical potential μ is not unusually large, the spectrum looks like a very narrow bow-tie shaped strip along the real axis, and the sets of eigenvalues with smallest absolute real parts and smallest magnitudes will nearly coincide. In the following we therefore define the deflation gap Δ as the largest deflated eigenvalue in magnitude, i.e., $\Delta = \max |\lambda_{\text{def}}|$.




Figure 1: Accuracy of the Krylov subspace approximation for $y = \text{sgn}(A)x$, where A is $\gamma_5 D_w(\mu)$ for a 6^4 lattice (for a lattice volume V the matrix $\gamma_5 D_w$ has dimension $12V$, such that $\dim(A) = 15552$ here). Left pane: Hermitian case ($\mu = 0$) using the Lanczos method, right pane: non-Hermitian case with chemical potential $\mu = 0.3$ using the two-sided Lanczos method. The relative error ε is shown as a function of the Krylov subspace size k for different deflation gaps Δ (given in parenthesis).

To analyze the efficiency of the algorithm we briefly sketch the three major contributions to the total CPU time. For each matrix A the deflation requires the computation of the critical eigenvalues and the corresponding eigenvectors. The time needed by the rest of the algorithm strongly depends on the eigenvalue gap, as the Krylov subspace size can be reduced if the deflation gap is increased. As mentioned at the beginning of this section, the product $f(A)x$ is usually needed for many source vectors x , e.g., as part of an iterative inversion. In this case the expensive deflation of A only needs to be performed once in an initialization step, while the Krylov subspace part of the algorithm will be repeated for each new vector x . For this reason we assume from now on that an initial deflation has been performed and we will concentrate on the efficiency of the Krylov subspace part of the algorithm. We discern two main components in the Krylov-Ritz method: the construction of the Krylov basis using the Lanczos or two-sided Lanczos algorithms, where the computation time grows linearly with the subspace size k , and the RHi to compute $\text{sgn}(H_k)$, which scales as k^3 . Figure 2 illustrates these last two contributions. For high accuracy the Krylov subspace becomes large such that the cost of the RHi dominates the total CPU time of the Krylov-Ritz approximation and the method becomes too costly. In the following, the implementation of the Krylov-Ritz approximation for which $\text{sgn}(H_k)$ is computed using Eq. (13) will be referred to as *non-nested method*. In the next section we will present a *nested* Krylov subspace method, which drastically reduces the cost to compute $\text{sgn}(H_k)e_1^{(k)}$ and vastly improves the overall efficiency of the Krylov-Ritz approximation.

4. Nested Krylov subspace method for the sign function

4.1. Nesting and preconditioning

We introduce a new method which speeds up the expensive computation of the vector $\text{sgn}(H_k)e_1^{(k)}$ required in the Krylov-Ritz approximations (10) and (12) to $\text{sgn}(A)x$. The idea is to approximate this matrix-vector product by a further Krylov-Ritz approximation, using a second, nested Krylov subspace (specified below) of size $\ell \ll k$, i.e.,

$$\text{sgn}(H_k)e_1^{(k)} = V_\ell \text{sgn}(H_\ell)e_1^{(\ell)}, \quad (14)$$

where V_ℓ is the matrix containing the basis vectors of the inner Krylov subspace, constructed with the Lanczos or two-sided Lanczos method, and H_ℓ is the inner Ritz or two-sided Ritz matrix. The $\text{sgn}(H_\ell)$ is computed using the RHi on the inner Ritz matrix H_ℓ . After substituting this result in Eq. (10) and (12) we get the nested approximation

$$y \approx |x| V_k V_\ell \text{sgn}(H_\ell)e_1^{(\ell)} \quad (15)$$




Figure 2: CPU time t (in seconds) versus accuracy for an 8^4 lattice configuration in the Hermitian case with deflation gap $\Delta = 0.055$ (left) and the non-Hermitian case with $\mu = 0.3$ and deflation gap $\Delta = 0.107$ (right). The full line shows the total time required to compute $\text{sgn}(A)x$, while the dashed line gives the time needed to construct the Krylov basis. The difference between both lines represents the time taken by the RHi to compute $\text{sgn}(H_k)$. The irregular convergence pattern for the non-Hermitian case is a well-known feature of the two-sided Lanczos algorithm.

to $\text{sgn}(A)x$. By introducing an additional Krylov subspace, the number of operations necessary to compute $\text{sgn}(H_k)e_1^{(k)}$ is reduced from $O(k^3)$ in the non-nested method to $O(\ell^3) + O(k\ell)$. If $\ell \ll k$ this will very much improve the efficiency of the Krylov-Ritz approximation.

The obvious choice for the inner Krylov subspace is $\mathcal{K}_\ell(H_k, e_1^{(k)})$. However, it is easy to see that approximations in this Krylov subspace will not improve the efficiency of the method. The Ritz matrix H_ℓ of the Krylov subspace $\mathcal{K}_\ell(H_k, e_1^{(k)})$ will only contain information coming from the $\ell \times \ell$ upper left corner of H_k , because of the tridiagonal nature of H_k and the sparseness of the source vector $e_1^{(k)}$. This will effectively cut down the size of the outer Krylov subspace from k to ℓ , which will substantially worsen the accuracy of the approximation if ℓ is chosen much smaller than k . Nonetheless, the nested Krylov subspace method can be made to work efficiently if we perform an initial preconditioning step on the tridiagonal Ritz matrix, replacing²

$$H_k \rightarrow H'_k = \frac{1}{2} [pH_k + (pH_k)^{-1}], \quad (16)$$

with p a positive real number, and construct the approximation to $\text{sgn}(H_k)e_1^{(k)}$ in the Krylov subspace $\mathcal{K}_\ell(H'_k, e_1^{(k)})$. This alternate Krylov subspace can be used to compute $\text{sgn}(H_k)e_1^{(k)}$ because the transformation leaves the sign unchanged. To show this, we note that both matrices have identical eigenvectors, as a matrix and its inverse share the same eigenvectors, and that the sign of their eigenvalues satisfies

$$\text{sgn} \frac{1}{2} \left(pz + \frac{1}{pz} \right) = \text{sgn} \text{Re} \left(pz + \frac{1}{pz} \right) = \text{sgn} \text{Re} \left(pz + \frac{pz^*}{|pz|^2} \right) = \text{sgn} \left[\underbrace{\left(1 + \frac{1}{|pz|^2} \right)}_{>0} \text{Re}(pz) \right] = \text{sgn}(z), \quad (17)$$

where we used the definition (6). Hence, $\text{sgn}(H'_k) = \text{sgn}(H_k)$ according to Eq. (4).³

As H_k is tridiagonal the cost of its inversion, required in (16), is only of $O(k)$. Moreover, as the transformation increases the relative gap between the spectrum and the singularity along the imaginary axis (see below), we expect a clear gain in efficiency for the inner Krylov-Ritz approximation, characterized by $\ell \ll k$.

For a Hermitian matrix the transformation induced by the preconditioning step is illustrated in Fig. 3 for real positive eigenvalues (for negative values the graph would be reflected with respect to the origin). The factor p is chosen to optimize the effect of the transformation on the relative distance to the imaginary axis, which in the Hermitian case

²The factor $1/2$ is chosen for convenience. For $p = 1$ the transformation actually mimics the first step of the RHi (13).

³If H_k is not diagonalizable, the equality can be shown by applying Eq. (17) to the integration variable in the integral representation (3).




Figure 3: Mapping of the preconditioning step $z' = (pz + 1/pz)/2$ for positive real eigenvalues and various values of p .

corresponds to a minimization of the condition number. We examine the condition number for the Hermitian case, assuming that the spectral support of H_k is similar to that of the original matrix A , after deflation. As can be seen from Fig. 3, after transformation the smallest eigenvalue (in absolute value) is $z'_{\min} = 1$, while the largest will be given by the transform of either the smallest or largest eigenvalues of H_k . The smallest condition number will be achieved when both values are identical, i.e., for p satisfying⁴

$$\frac{1}{2} \left(pz_{\min} + \frac{1}{pz_{\min}} \right) \stackrel{!}{=} \frac{1}{2} \left(pz_{\max} + \frac{1}{pz_{\max}} \right) \Rightarrow p_{\text{opt}} = \sqrt{\frac{1}{z_{\min} z_{\max}}}, \quad (18)$$

where $z_{\min} = \min |z|$ and $z_{\max} = \max |z|$, for z in the spectrum of H_k , and the largest transformed eigenvalue will be

$$z'_{\max} = \frac{1}{2} \left(\sqrt{\frac{z_{\max}}{z_{\min}}} + \sqrt{\frac{z_{\min}}{z_{\max}}} \right) \approx \frac{1}{2} \sqrt{\frac{z_{\max}}{z_{\min}}}. \quad (19)$$

In the Hermitian case, the transformation (16) therefore reduces the condition number C by a factor

$$\mathcal{F} = \frac{C}{C'} = \frac{z_{\max}}{z_{\min}} \left/ \frac{1}{2} \left(\sqrt{\frac{z_{\max}}{z_{\min}}} + \sqrt{\frac{z_{\min}}{z_{\max}}} \right) \right. \approx 2 \sqrt{\frac{z_{\max}}{z_{\min}}}. \quad (20)$$

The effect of the preconditioning of the Ritz matrix for a typical spectrum of $\gamma_5 D_w$ in lattice QCD is illustrated in Fig. 4 for the Hermitian case. The top and bottom graphs depict the spectra of H_k and H'_k , respectively. The spectrum of the original Ritz matrix has only a small gap at zero, while the gap for the transformed matrix is large. In this example, the condition number is almost improved by a factor 20. In general, the value of z_{\max} for $\gamma_5 D_w$ varies only slightly with the choice of the simulation parameters and \mathcal{F} will mainly depend on the deflation gap.

For the non-Hermitian case, let us assume that the complex spectrum is contained in the circles $C(-m, r) \cup C(m, r)$, with real center $m > 0$ and radius $r < m$. The optimal p , maximizing the relative distance from the imaginary axis for the transformed spectrum, is still given by Eq. (18) which now simplifies to $p_{\text{opt}} = (m^2 - r^2)^{-1/2}$. For this choice the transformed eigenvalues are contained in the circles $C(-m', r') \cup C(m', r')$ with center $m' = (m_s + 1/m_s)/2$ and radius $r' = (m_s - 1/m_s)/2$, where $m_s \equiv p_{\text{opt}} m$. This is illustrated in the left panel of Fig. 5, where we show the transformation of a circle $C(m, r)$ for the optimal and a sub-optimal value of p . For sub-optimal p the transformation yields an inner and an outer circle-like contour, which merge into the circle $C(m', r')$ when $p \rightarrow p_{\text{opt}}$. For p_{opt} the relative distance from the imaginary axis will be maximal and we expect the transformation (16) to work best. The gain in efficiency will however not be as large as for the Hermitian case. This can be quantified by the relative distance to the imaginary

⁴In practice p_{opt} is only known approximately, as it is computed from spectral information of A instead of H_k . However, this has no significant impact on the performance of the nested method.




Figure 4: Upper pane: density of eigenvalues of H_k in the Hermitian case for an 8^4 lattice with $k = 1536$. The spectrum has a narrow deflation gap $\Delta = 0.055$. The optimal p -factor (18) for the transformation (16) is $p_{\text{opt}} \approx 1.86$ (using $z_{\text{min}} = \Delta$ and $z_{\text{max}} = 5.26$). The lower pane shows the corresponding eigenvalue density of the transformed matrix H'_k , where the condition number is improved by a factor $\mathcal{F} = 19.3$ (see Eq. (20)).

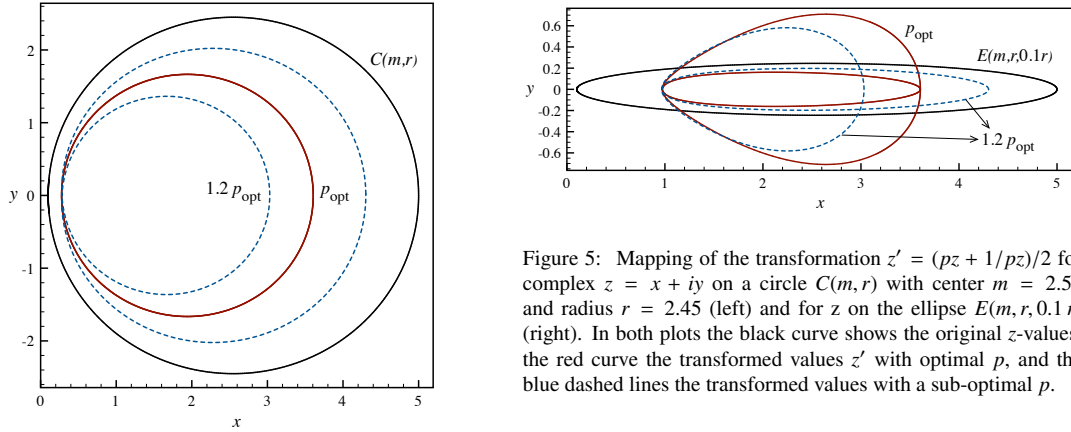


Figure 5: Mapping of the transformation $z' = (pz + 1/pz)/2$ for complex $z = x + iy$ on a circle $C(m, r)$ with center $m = 2.55$ and radius $r = 2.45$ (left) and for z on the ellipse $E(m, r, 0.1r)$ (right). In both plots the black curve shows the original z -values, the red curve the transformed values z' with optimal p , and the blue dashed lines the transformed values with a sub-optimal p .

axis, in analogy to the calculation performed above for the Hermitian case. For the original spectrum we define the relative distance as

$$d \equiv \frac{\min |\operatorname{Re} z|}{\max |\operatorname{Re} z|} = \frac{m - r}{m + r} \quad (21)$$

and for the transformed spectrum

$$d' \equiv \frac{\min |\operatorname{Re} z'|}{\max |\operatorname{Re} z'|} = \frac{m' - r'}{m' + r'} = \frac{1}{m_s^2} = \frac{m^2 - r^2}{m^2}. \quad (22)$$

The improvement factor due to the transformation is given by the ratio of these distances, yielding

$$\mathcal{F} = \frac{d'}{d} = \left(\frac{m + r}{m} \right)^2 = \left(2 - \frac{\Delta}{m} \right)^2, \quad (23)$$

where we wrote $r = m - \Delta$, with Δ the deflation gap. When $\Delta \ll m$ we will have $\mathcal{F} \approx 4$. For the example shown in the left plot of Fig. 5 the transformation generates an improvement by a factor $\mathcal{F} = 3.84$, as computed with Eq. (23).




Figure 6: Red dots: spectrum of H_k in the non-Hermitian case for an 8^4 lattice with $k = 1580$, $\mu = 0.3$ and deflation gap $\Delta = 0.107$. The optimal p -factor (18) for the transformation (16) is $p_{\text{opt}} \approx 1.335$ (using $z_{\text{min}} = \Delta$ and $z_{\text{max}} = 5.243$). Blue circles: the corresponding spectrum of the transformed matrix H'_k . As desired, the transformed eigenvalues are well away from the imaginary axis (the vertical lines at $x = \pm 1$ serve to guide the eye). Note the different scales on the x and y axes.

In lattice QCD at nonzero baryon density $\gamma_5 D_w$ is usually weakly non-Hermitian and, after deflation, the spectra are contained in ellipses $E(-m, a, b) \cup E(+m, a, b)$, with center $m \in \mathbb{R}^+$ and major and minor axes a and b along the real and imaginary axes, respectively. The transformation (16) of an ellipse $E(m, a, b)$ with aspect ratio $a/b = 10$ is illustrated in the right panel of Fig. 5. For such a narrow ellipse the transformed spectrum is qualitatively similar to the Hermitian case, as all the eigenvalues are transformed to the right of $z' = 1$, i.e., away from the imaginary axis, such that the high efficiency of the transformation is still guaranteed. The optimal value p_{opt} is again determined by (18) with $p_{\text{opt}} = (m^2 - a^2)^{-1/2}$, as it maximizes the relative distance from the imaginary axis. The transformation is illustrated for a realistic test case of lattice QCD in Fig. 6, where the eigenvalues and transformed eigenvalues of the Ritz matrix for $\gamma_5 D_w(\mu)$ are shown for $\mu = 0.3$.

As we will see below the preconditioning step significantly speeds up the Krylov-Ritz approximation in its application to lattice QCD at zero and nonzero chemical potential.

4.2. Convergence

In this section we investigate the convergence properties of the nested method. The method was implemented to compute the sign function of $\gamma_5 D_w(\mu)$ needed by the overlap operator (1), for both the Hermitian and the non-Hermitian case. Whenever the matrix has eigenvalues close to the imaginary axis, these critical eigenvalues are first deflated to open up a deflation gap, necessary to keep the Krylov subspace within a reasonable size (see Sec. (3)). Our implementation uses Chroma [19] to compute the Wilson operator. The linear algebra is performed with BLAS and LAPACK routines. To ensure the efficiency of the nested method a judicious implementation of the preconditioning step (16), used to construct the inner Krylov subspace, is needed. Explicitly inverting the tridiagonal matrix H_k to form the full matrix H'_k , then constructing the basis of the inner Krylov subspace by successive full matrix-vector multiplications would make a rather inefficient algorithm. To construct the inner Krylov subspace we do not need to construct the full matrix H'_k explicitly, but only have to apply H'_k to $\ell - 1$ vectors of \mathbb{C}^k (in the non-Hermitian case $H_k'^\dagger$ is also needed). These products are best computed using the LU decomposition of H_k , which is $O(k)$ and thus especially efficient for tridiagonal matrices. A detailed listing of the algorithm is given in Appendix C.

The overall accuracy of the nested approximation (15) depends on the parameters k and ℓ , defining the sizes of the outer and inner Krylov subspaces, respectively. For $\ell \rightarrow k$ the solution of the nested method will converge to that of the non-nested method with Krylov subspace size k and accuracy ε_k , so its total error will also converge to ε_k . To investigate the accuracy of the nested algorithm, our strategy is to fix the outer Krylov subspace size k , corresponding to a certain desired accuracy, and vary the inner Krylov subspace size ℓ . We show the convergence results for an 8^4 lattice configuration in Fig. 7, for both the Hermitian and non-Hermitian case. As expected the nested method reaches the accuracy of the non-nested method when its size is large enough. Surprisingly however, this happens for $\ell \ll k$,




Figure 7: Accuracy ε of the nested method for an 8^4 lattice configuration. Hermitian case with deflation gap $\Delta = 0.055$ (left) and non-Hermitian case with $\mu = 0.3$ and deflation gap $\Delta = 0.107$ (right). $\varepsilon(k)$ shows how the error of the non-nested method decreases with growing Krylov subspace (blue line). The vertical line fixes the size k of the outer Krylov space used in the nested method. $\varepsilon(\ell)$ shows the accuracy of the nested method, for fixed k , as a function of the size ℓ of the inner Krylov subspace (red line). The rapid convergence illustrates the efficiency of the nested method. The smallest value of ℓ for which optimal convergence is reached is denoted by ℓ_{opt} . Note that we always restrict ourselves to even Krylov subspace sizes, as odd values systematically give a somewhat worse accuracy because of spurious near-zero eigenvalues occurring in the Ritz matrix.




Figure 8: Convergence of the nested method for an 8^4 lattice configuration as a function of the relative inner Krylov subspace size ℓ/k , for various deflation gaps (given in parenthesis). For each gap the value of k is chosen such that an accuracy of 10^{-8} is achieved. Left: Hermitian case with $k = 2806, 1462$ and 758 for deflation gap $\Delta = 0.025, 0.05$ and 0.1 . Right: non-Hermitian case with $\mu = 0.3$ and $k = 3808, 1456$ and 634 for $\Delta = 0.05, 0.1$ and 0.2 . Again, the irregular convergence pattern for the non-Hermitian case is characteristic for the two-sided Lanczos algorithm.

as the convergence of the inner Ritz approximation seems to be extremely fast. The smallest value of ℓ for which optimal convergence is reached will be called ℓ_{opt} . The fast convergence is closely related to the large improvement in condition number discussed in the previous section. We also showed in Eq. (20) how the improvement of the condition number, due to the preconditioning of the Ritz matrix H_k , depends on the deflation gap. A smaller gap will yield a larger improvement, and vice-versa. This in turn will influence the convergence rate of the nested method. Figure 8 verifies that the result $\ell_{\text{opt}} \ll k$ remains valid for different deflation gaps. The figure also illustrates that the somewhat larger reduction in condition number achieved for a smaller gap yields an accordingly smaller ratio ℓ_{opt}/k (approximately proportional to the ratio of the respective improvement factors \mathcal{F}). This is an additional advantage as the size reduction is largest when the outer subspace is large. In all cases, the inner Krylov subspace can be taken much smaller than the outer subspace, such that the efficiency of the Krylov-Ritz method is substantially boosted, as will be shown in the benchmarks below.




Figure 9: Error and CPU usage of the nested method for lattice size 8^4 . Hermitian case with deflation gap $\Delta = 0.055$ and fixed $k = 1536$ (left), and non-Hermitian case with $\mu = 0.3$, $\Delta = 0.107$ and $k = 1580$ (right). $\varepsilon(\ell)$ shows the accuracy versus inner Krylov subspace size ℓ (red line). $t(\ell)$ shows the total CPU time in seconds (solid blue line), while the horizontal dashed line measures the time needed to construct the basis in the outer Krylov subspace. The difference between both lines corresponds to the time taken by the RHi to compute $\text{sgn}(H_\ell)$. The vertical band highlights the operational window of the nested method, i.e., the region in ℓ where the accuracy is optimal, but the CPU-time used to compute $\text{sgn}(H_\ell)$ is negligible.

We also verified that the convergence curves are fairly insensitive to the choice of the source vector and lattice configuration. The fast convergence property of the nested method is generic, regardless of the simulation details, for both the Hermitian and non-Hermitian case, even though the precise value of ℓ_{opt} depends on the lattice size, the simulation parameters, the deflation gap and the desired overall accuracy (determined by k).

4.3. Benchmarks

With the fast convergence ($\ell_{\text{opt}} \ll k$) discussed in the previous section, we can expect a substantial gain in computation time when using the nested method. The total CPU time consumed by the nested method is illustrated in Fig. 9 for the Hermitian case (left) and the non-Hermitian case (right). The size of the outer Krylov subspace is kept fixed, such that its construction gives a constant contribution to the run time, depicted by the horizontal dashed line. The contribution to the CPU time which varies with ℓ mainly comes from the computation of $\text{sgn}(H_\ell)$ with the RHi and is proportional to ℓ^3 . For $\ell \approx k$ the total run time of the nested method is about equal to that of the non-nested method. However, as illustrated by the $\varepsilon(\ell)$ curve (red line) and discussed in Sec. 4.2, ℓ can be chosen much smaller while preserving the accuracy of the non-nested method. The central result, illustrated by the vertical band in Fig. 9, is that there exists an interval in ℓ for which the accuracy is still optimal, but the CPU time needed to compute $\text{sgn}(H_\ell)$ with the RHi is negligible compared to the time required to construct the outer Krylov subspace. There is therefore no need to make a compromise between run time and accuracy, as both can be optimized simultaneously. The error in this range is the minimal error achievable with the given size of the outer Krylov subspace, while the run time is completely dominated by the cost for building the basis in that subspace. The nested method is able to quench the CPU time needed for the computation of $\text{sgn}(H_k)e_1^{(k)}$ without affecting the accuracy of the Krylov-Ritz approximation.

To evaluate the nested method further, we compare it to state-of-the-art rational approximation methods. In the Hermitian case the Zolotarev rational approximation, evaluated with a multi-shift conjugate gradient inverter [6], is routinely used in lattice simulations. In the non-Hermitian case, i.e., simulations at nonzero baryon density, overlap fermions are not yet commonly used because of their high cost, but recently an efficient algorithm was presented, which evaluates the Neuberger rational approximation using a multi-shift restarted FOM inverter [10]. In Fig. 10 we compare the results obtained with the nested Krylov subspace and rational approximation methods, and show how the CPU time varies as a function of the achieved accuracy for various lattice sizes. In all cases the Hermitian and non-Hermitian versions of the nested method perform better than the rational approximation method. The volume dependence of the run time for a fixed accuracy ε can be extracted from Fig. 10 and is displayed for $\varepsilon = 10^{-8}$ in Fig. 11. Fits to the nested method results show a volume dependence which is slightly steeper than linear, i.e., proportional to $V^{1.2}$ for the Hermitian case and $V^{1.3}$ for the non-Hermitian case. The comparisons clearly demonstrate the good efficiency of the nested method.

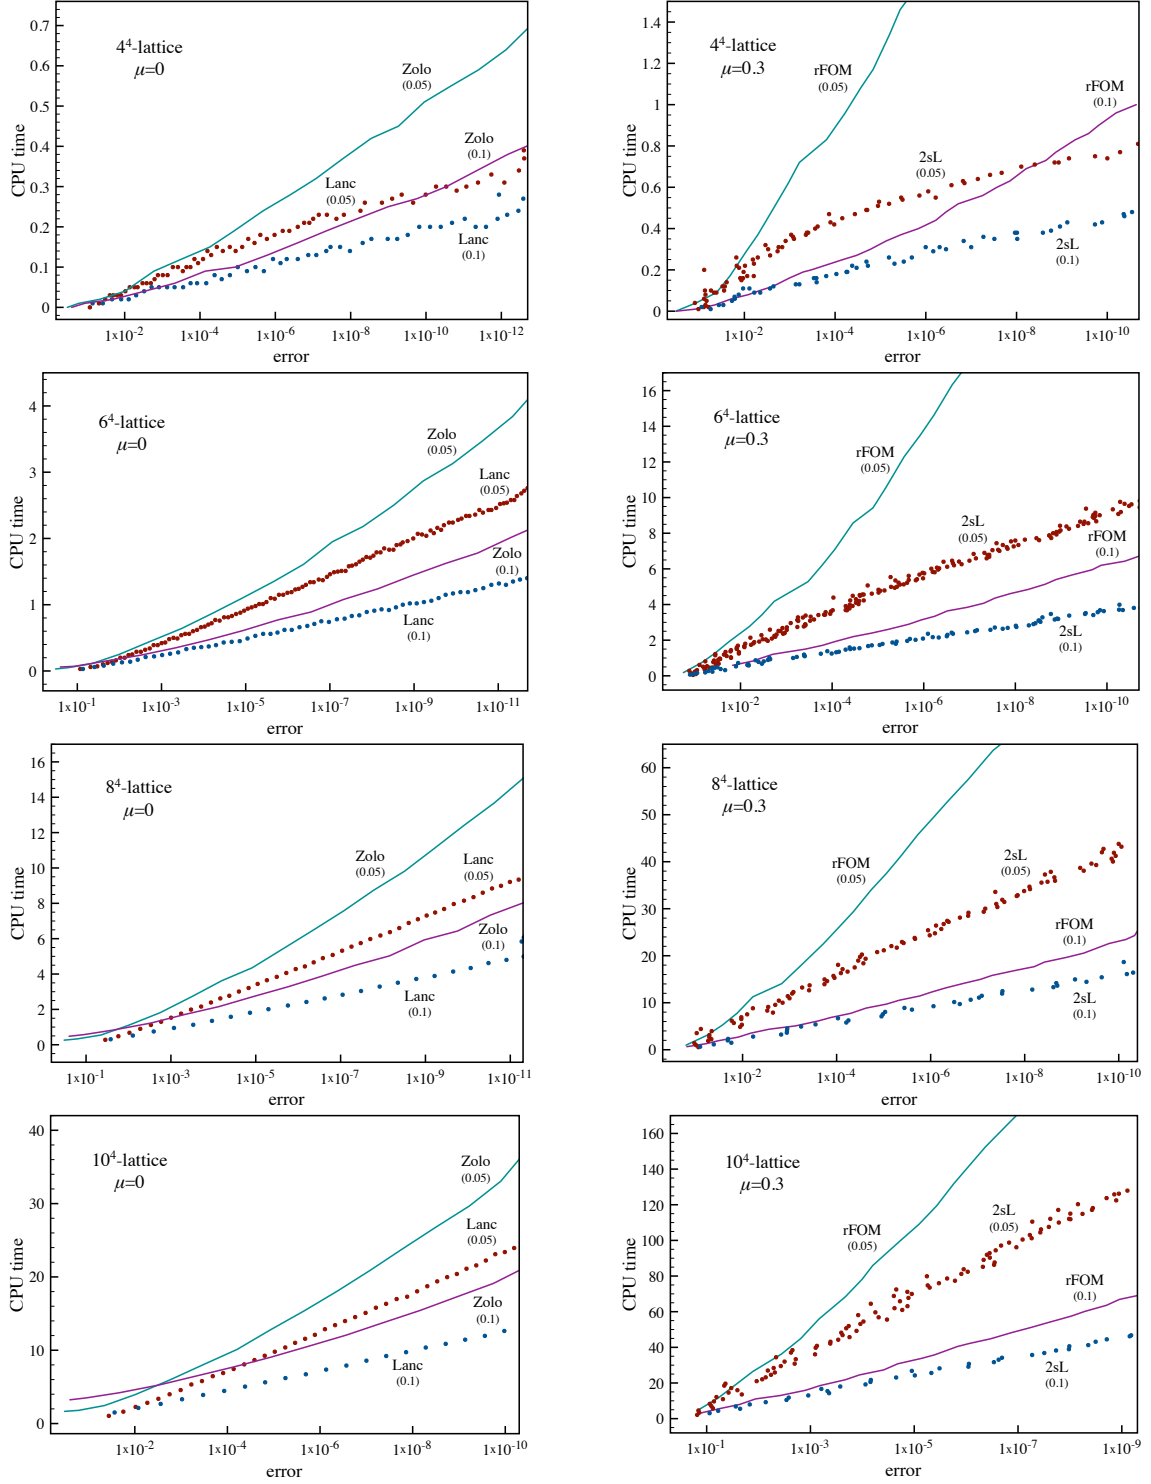


Figure 10: Comparison of the nested Krylov subspace method (filled circles) with rational approximation methods (full lines) for lattices of sizes 4^4 , 6^4 , 8^4 and 10^4 for two different deflation gaps (given in parenthesis). Left: Hermitian case comparing the nested Lanczos approximation (Lanc) with the Zolotarev approximation (Zolo), evaluated using the Chroma QCD library. Right: non-Hermitian case with $\mu = 0.3$ comparing the nested two-sided Lanczos method (2sL) with the Neuberger approximation evaluated with a restarted FOM algorithm (rFOM). The timings were measured on a single 2.4 GHz Intel Core 2 core with 8 GB of memory.




Figure 11: Volume dependence of the run times for the nested method and rational approximation methods for the Hermitian (left) and non-Hermitian (right) case. The data are taken from Fig. 10 at an accuracy of $\varepsilon = 10^{-8}$.

4.4. Note on the memory usage

In the numerical tests we observed that, for a fixed deflation gap, the Krylov subspace size needed to achieve a certain accuracy is almost independent of the lattice volume in the Lanczos approximation and only grows slowly with the volume in the two-sided Lanczos approximation. Therefore, the memory consumed by the Krylov basis V_k is roughly proportional to the lattice volume. For large lattice sizes this storage requirement might become too large to run the Krylov-Ritz approximation on a single node.

One solution, which only requires little storage, is to implement a *double-pass* version of the algorithm, which is possible due to the use of short recurrences. In double-pass mode only the two most recently generated basis vectors are stored during the construction of the outer Krylov subspace basis. In the first pass the matrix H_k is built and the product $\text{sgn}(H_k)e_1^{(k)}$ is computed with Eq. (14). In the second pass the basis vectors of the outer Krylov subspace are generated again and immediately added in a linear combination, whose coefficients were computed in the first pass. The drawback of this variant is that the Krylov basis is constructed twice, such that the corresponding CPU time will be doubled.

The more efficient solution is to parallelize the single-pass version of the algorithm, such that the memory requirement gets distributed over the available nodes. Benchmarks on larger volumes, using such a parallel implementation, are currently being performed.

4.5. Multi-level nesting

In principle, if the inner Krylov subspace in Eq. (15) is still too large for an efficient application of the RHi on the inner Ritz matrix, the nested method could be applied recursively.⁵ In this case we rename k to k_0 , ℓ to k_1 , and add more recursively nested levels k_i as necessary. Except for the deepest level, the matrix-vector product $\text{sgn}(H_{k_i})e_1^{(k_i)}$ required at level i will be computed with a Krylov-Ritz approximation (14) in the nested Krylov subspace $\mathcal{K}_{k_{i+1}}(H'_{k_i}, e_1^{(k_i)})$, where H'_{k_i} is defined by Eq. (16) on H_{k_i} and typically $k_{i+1} \ll k_i$. At the deepest level the sign function of the Ritz matrix will be evaluated with the RHi. This multi-level nesting is illustrated in Fig. 12, where we show the convergence curves for 1, 2, 3, 4 and 5 nested levels as a function of the size of the innermost Krylov subspace, with the sizes of all outer levels kept fixed to some value inside their convergence region (the convergence curves do not depend on the precise choice of the outer k_i 's). As before, the convergence criterion is set by the size k_0 of the outer Krylov subspace. Each additional level lowers the size of the Krylov subspace. In the case depicted in Fig. 12 the optimal Krylov subspace sizes, i.e. where convergence is reached, for the successive levels decreases from $1536 \rightarrow 90 \rightarrow 20 \rightarrow 8 \rightarrow 4 \rightarrow 2$. The improvement is most dramatic for the first nested level, but fast convergence is exhibited at all levels⁶. This can

⁵Note that for all cases considered in the current study a single level of nesting was sufficient.

⁶For each level the p -factor for Eq. (16) is computed using Eq. (18), using appropriately approximated boundaries for the spectrum of the Ritz matrix of the previous level. Note that the factor p converges to 1 as more levels are introduced, and the preconditioning step converges to the RHi.




Figure 12: Accuracy ε of the nested method with 1, 2, 3, 4 and 5 nesting levels for lattice size 8^4 in the Hermitian case with deflation gap $\Delta = 0.055$ and $k_0 = 1536$. We plot the dependence of ε on the size k_i , $i = 1, \dots, 5$, of the innermost Krylov space. The convergence curves are labelled with the number of levels in the method. For the i -level method, the outer levels k_j , $j = 1, \dots, i - 1$ are fixed to a value in their convergence region.

be related to the quadratically convergent RHi, as the preconditioning step at each level mimics a step of the RHi and compresses the spectrum more and more towards ± 1 . Moreover, the judicious choice of p at each level improves the convergence even more. It is intriguing to note that, in the example of Fig. 12, the sign of a matrix of dimension $n = 49152$ can be evaluated to an accuracy of 10^{-9} by computing the sign of a 2×2 matrix, which is then lifted back to the original n -dimensional space through linear combinations of Krylov vectors. This emphasizes again the power of Krylov subspace methods.

5. Conclusions

In this paper we have presented a nested Krylov subspace method which boosts the Krylov-Ritz approximations used to compute the sign function of both Hermitian and non-Hermitian matrices. The Krylov-Ritz approximation projects the matrix on a Krylov subspace in which it computes the sign function exactly, before lifting it back to the original space. Its standard implementation suffers from the CPU intensive computation of the sign of the Ritz matrix, which goes like the cube of the Krylov subspace size. By making an additional projection on a much smaller Krylov subspace, the nested method significantly reduces the total computation time of the Krylov-Ritz approximation, without affecting its accuracy. Numerical tests showed that the nested method works equally well for Hermitian and non-Hermitian matrices and is more efficient than state-of-the-art rational approximation methods. Moreover, it exhibits a good, close to linear, volume scaling. We are currently investigating the efficiency of the nested method for larger lattice volumes using a parallel implementation of the algorithm.

To end, we comment on the relation between the nested method and the extended Krylov subspace methods introduced in Ref. [20]. An extended Krylov space is defined as

$$\mathcal{K}_k(A, A^{-1}, x) = \text{span}(x, Ax, A^{-1}x, A^2x, A^{-2}x, \dots, A^{k-1}x, A^{-k+1}x), \quad (24)$$

and an approximation in that subspace approximates $f(A)$ by the sum $Q(A) = \sum_{i=-k+1}^{k-1} c_i A^i$. In the nested method we construct the ℓ -dimensional Krylov subspace $\mathcal{K}_\ell(H'_k, e_1^{(k)})$, which forms an ℓ -dimensional subspace of the $(2\ell - 1)$ -dimensional extended Krylov subspace $\mathcal{K}_\ell(pH_k, (pH_k)^{-1}, e_1^{(k)})$. The nested method implicitly fixes the coefficients of the positive and negative powers of $Q(H_k)$ to be equal, $c_{-i} = c_i$, which follows from the use of the property $\text{sgn}(H_k) = \text{sgn}(H_k + H_k^{-1})$. Hence, the nested method implicitly truncates the size of the extended Krylov subspace.

Approximations for the sign function in extended Krylov subspaces have been briefly considered recently [21], however not in combination with the nesting of Krylov subspaces, i.e. the extended subspace is constructed for the original matrix A , not for the Ritz matrix H_k . Evidently this is not feasible in the application to lattice QCD as the inversion of the γ_5 -Wilson Dirac operator is too expensive in order to construct extended Krylov subspaces.

To conclude, we briefly consider the application of the nested method to other matrix functions. The method presented in Sec. 4 requires a transformation which leaves the matrix function invariant, similar to Eq. (16) for the sign function. If such a transformation is not known, the nested method could be adapted by using an extended Krylov subspace method at the inner level. This is also a topic of work in progress.

Acknowledgements

We would like to thank Andreas Frommer and Tilo Wettig for useful discussions.

Appendix A. Lanczos algorithm

```

 $v_1 \leftarrow \frac{x}{|x|}$ 
 $r \leftarrow Av_1$ 
for  $j = 1$  to  $k$  do
   $H(j, j) \leftarrow v_j^\dagger r$ 
   $r \leftarrow r - H(j, j)v_j$ 
  if  $j = k$  then
    stop
  end if
   $\beta \leftarrow \sqrt{r^\dagger r}$ 
   $H(j, j+1) \leftarrow \beta$ 
   $H(j+1, j) \leftarrow \beta$ 
   $v_{j+1} \leftarrow \frac{1}{\beta}r$ 
   $r \leftarrow Av_{j+1}$ 
   $r \leftarrow r - \beta v_j$ 
end for

```

All $H(i, j)$ not assigned above are zero. Consequently H is tridiagonal and symmetric. The v_j are the column vectors of the matrix V_k .

Appendix B. Two-sided Lanczos algorithm

```

 $v_1 \leftarrow \frac{x}{|x|}$ 
 $w_1 \leftarrow v_1$ 
 $r \leftarrow Av_1$ 
 $l \leftarrow A^\dagger w_1$ 
for  $j = 1$  to  $k$  do
   $H(j, j) \leftarrow w_j^\dagger r$ 
   $r \leftarrow r - H(j, j)v_j$ 
   $l \leftarrow l - (H(j, j))^* w_j$ 
  if  $j = k$  then
    stop
  end if
   $\delta \leftarrow r^\dagger l$ 
  if  $\delta = 0$  then
    serious breakdown, stop
  end if
   $\beta \leftarrow \sqrt{\delta}$ 
   $H(j+1, j) \leftarrow \beta$ 
   $\gamma \leftarrow \frac{\delta^*}{\beta}$ 
   $H(j, j+1) \leftarrow \gamma$ 
   $v_{j+1} \leftarrow \frac{1}{\beta}r$ 

```

$$\begin{aligned}
w_{j+1} &\leftarrow \frac{1}{\gamma^*} l \\
r &\leftarrow A v_{j+1} \\
l &\leftarrow A^\dagger w_{j+1} \\
r &\leftarrow r - \gamma v_j \\
l &\leftarrow l - \beta^* w_j
\end{aligned}$$

end for

The v_j and w_j are the column vectors of the matrices V_k and W_k , respectively. All $H(i, j)$ not assigned above are zero. Consequently H is tridiagonal, but not symmetric as in the Hermitian case. The coefficients β and γ are, non-uniquely, chosen to satisfy the biorthonormality condition

$$w_j^\dagger v_i = \delta_{ij}. \quad (\text{B.1})$$

There are potential problems in the two-sided Lanczos process, namely serious breakdowns and near breakdowns, where $\delta \leftarrow r^\dagger l = 0$, respectively ≈ 0 , however, these were not encountered in our numerical tests.

Appendix C. Nested algorithm

Given a (non-)Hermitian matrix A , a source vector x and the critical eigenvectors r_i (left and right eigenvectors l_i and r_i), with eigenvalues λ_i , $i = 1, \dots, m$, do:

1. Apply Left-Right deflation (see Ref.[11]) to construct x_\ominus , where the components of the source vector x along the eigenvectors r_i have been removed:

$$x_\ominus = x - \sum_{i=1}^m \langle l_i, x \rangle r_i,$$

where $l_i = r_i$ for Hermitian A .

2. Run the (two-sided) Lanczos algorithm from Appendix A (Appendix B) with A and x_\ominus to obtain V_k and H_k .
3. Perform an LU decomposition of pH_k , e.g., with the LAPACK routine `dgtrf` (`zgtrf`). This yields a lower triangular matrix L with unit diagonal and one sub-diagonal, and an upper triangular matrix U with one diagonal and two super-diagonals. All other entries of L and U are zero.
4. Run the (two-sided) Lanczos algorithm with $H'_k = (pH_k + (pH_k)^{-1})/2$ and source vector $e_1^{(k)}$ to construct the Krylov basis V_ℓ and the Ritz matrix H_ℓ . To do so apply H'_k to each Krylov vector v :
 - (a) Compute $(pH_k)^{-1}v$ using a sparse LU back substitution, e.g., with the LAPACK routine `dgtrfs` (`zgtrfs`).
 - (b) Compute $(pH_k)v$ and add to the result of (a). This tridiagonal multiply and add can be done efficiently using the BLAS band-matrix-vector multiplication routine `dsbmv` (`zgbmv`).
5. Run the RHi (or any other suitable method to compute the sign function) on H_ℓ to obtain $\text{sgn}(H_\ell)$.
6. The final approximation is then given by

$$\text{sgn}(A)x \approx \sum_{i=1}^m \text{sgn}(\lambda_i) \langle l_i, x \rangle r_i + |x_\ominus| V_k V_\ell \text{sgn}(H_\ell) e_1^{(\ell)}.$$

Note that steps (3-5) are done in real arithmetic in the Hermitian case.

References

- [1] R. Narayanan and H. Neuberger, *Chiral determinant as an overlap of two vacua*, *Nucl. Phys.* **B412** (1994) 574–606, [[hep-lat/9307006](#)].
- [2] R. Narayanan and H. Neuberger, *A construction of lattice chiral gauge theories*, *Nucl. Phys.* **B443** (1995) 305–385, [[hep-th/9411108](#)].
- [3] J. C. R. Bloch and T. Wettig, *Overlap Dirac operator at nonzero chemical potential and random matrix theory*, *Phys. Rev. Lett.* **97** (2006) 012003, [[hep-lat/0604020](#)].
- [4] J. C. R. Bloch and T. Wettig, *Domain-wall and overlap fermions at nonzero quark chemical potential*, *Phys. Rev.* **D76** (2007) 114511, [[arXiv:0709.4630](#)].
- [5] H. Neuberger, *A practical implementation of the overlap-Dirac operator*, *Phys. Rev. Lett.* **81** (1998) 4060–4062, [[hep-lat/9806025](#)].

- [6] J. van den Eshof, A. Frommer, T. Lippert, K. Schilling, and H. A. van der Vorst, *Numerical methods for the QCD overlap operator. I: Sign-function and error bounds*, *Comput. Phys. Commun.* **146** (2002) 203–224, [[hep-lat/0202025](#)].
- [7] T.-W. Chiu, T.-H. Hsieh, C.-H. Huang, and T.-R. Huang, *A note on the Zolotarev optimal rational approximation for the overlap Dirac operator*, *Phys. Rev.* **D66** (2002) 114502, [[hep-lat/0206007](#)].
- [8] A. D. Kennedy, *Fast evaluation of Zolotarev coefficients*, [hep-lat/0402038](#).
- [9] E. I. Zolotarev, *Application of elliptic functions to the question of functions deviating least and most from zero*, *Zap. Imp. Akad. Nauk. St. Petersburg* **30** (1877) 5.
- [10] J. C. R. Bloch, T. Breu, A. Frommer, S. Heybrock, K. Schäfer, and T. Wettig, *Short-recurrence Krylov subspace methods for the overlap Dirac operator at nonzero chemical potential*, *Comput. Phys. Commun.* **181** (2010) 1378–1387, [[arXiv:0910.1048](#)].
- [11] J. C. R. Bloch, A. Frommer, B. Lang, and T. Wettig, *An iterative method to compute the sign function of a non-Hermitian matrix and its application to the overlap Dirac operator at nonzero chemical potential*, *Comput. Phys. Commun.* **177** (2007) 933–943, [[arXiv:0704.3486](#)].
- [12] J. C. R. Bloch, T. Breu, and T. Wettig, *Comparing iterative methods to compute the overlap Dirac operator at nonzero chemical potential*, *PoS LATTICE2008* (2008) 027, [[arXiv:0810.4228](#)].
- [13] H. Neuberger, *Exactly massless quarks on the lattice*, *Phys. Lett.* **B417** (1998) 141–144, [[hep-lat/9707022](#)].
- [14] P. Hasenfratz and F. Karsch, *Chemical potential on the lattice*, *Phys. Lett.* **B125** (1983) 308.
- [15] G. Golub and C. V. Loan, *Matrix Computations*. The John Hopkins University Press, 1989.
- [16] E. Gallopoulos and Y. Saad, *On the parallel solution of parabolic equations*, in *Proceedings of the International Conference on Supercomputing 1989, Heraklion, Crete, June 5-9, 1989* (R. D. Groot, ed.), ACM press, 1989.
- [17] Y. Saad, *Analysis of some Krylov subspace approximations to the matrix exponential operator*, *SIAM J. Numer. Anal.* **29** (1992), no. 1 209–228.
- [18] J. Roberts, *Linear model reduction and solution of the algebraic Riccati equation by use of the sign functions*, *Internat. J. Control* **32** (1980) 677–687.
- [19] SciDAC Collaboration, R. G. Edwards and B. Joo, *The Chroma software system for lattice QCD*, *Nucl. Phys. Proc. Suppl.* **140** (2005) 832, [[hep-lat/0409003](#)].
- [20] V. Druskin and L. Knizhnerman, *Extended Krylov subspaces: Approximation of the matrix square root and related functions*, *SIAM J. Matrix Anal. Appl.* **19** (1998) 755–771.
- [21] L. Knizhnerman and V. Simoncini, *A new investigation of the extended Krylov subspace method for matrix function evaluations*, *Numer. Linear Algebra Appl.* **17** (2010), no. 4 615–638.